

**Inequalities in Child Development in Peru:
Evidence about its Origins and the Effects of Policy
Interventions on Parental Behaviour**

Thesis submitted in partial fulfilment of the requirements
for the Degree of Doctor of Philosophy

by

Juan Francisco Castro

St Cross College

Department of International Development
University of Oxford

MT 2015

DECLARATION OF AUTHORSHIP

Name: JUAN FRANCISCO CASTRO
College: ST CROSS

Candidate number: 628518
Supervisor(s) of thesis/essay:
Professor Stefan Dercon
Dr. Caine Rolleston

Title of thesis:
Inequalities in Child Development in Peru: Evidence about its Origins and the Effects of Policy Interventions on Parental Behaviour

Word count: 72,055

There is extensive information and guidance on academic good practice and plagiarism on the University website: <www.admin.ox.ac.uk/epsc/plagiarism>. You must read all information contained here, and complete the on-line plagiarism course.

Please tick to confirm the following:

I am aware of the University's disciplinary regulations concerning conduct in examinations and, in particular, of the regulations on plagiarism (c.f. The Proctors' and Assessor's Memorandum, Section 9.6 at www.admin.ox.ac.uk/proctors/info/pam/section9.shtml). ☒

The thesis I am submitting is entirely my own work except where otherwise indicated. ☒

It has not been submitted, either wholly or substantially, for another degree of this University, or for a degree at any other institution. ☒

I have clearly signalled the presence of quoted or paraphrased material and referenced all sources. ☒

I have acknowledged appropriately any assistance I have received in addition to that provided by my supervisor(s). ☒

I have not sought assistance from any professional agency. ☒

I have not repeated any material from other pieces of work that I have previously submitted for assessment for this degree, except where permitted. ☒

I agree to retain an electronic copy of this work until the publication of my final examination result, except where submission in hand-written format is permitted. I agree to make any such electronic copy available to the examiners should it be necessary to confirm my word count or to check for plagiarism.

I am the sole author of all the chapters of this thesis except the second. The second chapter is based on joint research with Caine Rolleston, although the document in this thesis is largely my own work.

Candidate's signature:



Date: 11 / January / 2016

Inequalities in Child Development in Peru: Evidence about its Origins and the Effects of Policy Interventions on Parental Behaviour

by

Juan Francisco Castro

St Cross College

Department of International Development
University of Oxford

Thesis submitted in partial fulfilment of the requirements
for the Degree of Doctor of Philosophy

MT 2015

Abstract

This thesis examines two topics related to the prevalence of unequal developmental outcomes among children of different socioeconomic backgrounds in Peru. The first topic focuses on the origins of these inequalities by measuring the relative importance of school and early childhood influences for the emergence of cognitive development gaps between urban and rural children. I use rich longitudinal information on cognitive achievement, household and school characteristics, and a novel decomposition strategy. Results reveal that observable school influences occurring between ages 6 and 8, account for a significant share (around 35% and no less than 28%) of the difference in cognitive skill. The share attributable to differences in the early childhood environment is important but no larger than 50%. I also discuss how the proposed decomposition strategy compares to other linear decomposition methods and why it is less prone to biases than those employed so far in the literature. The second topic is related to the effects and mechanisms behind early childhood development interventions that seek to improve the home environment through a change in parental behaviour. I use data on parenting practices and caregiver beliefs regarding the importance of parent-child interactions collected from the control and treatment groups of families randomized for the evaluation of a home-visiting programme recently launched at scale in rural Peru. I find that this intervention produced a positive effect on the quality of the home environment ($d = 0.5$; $p < 0.01$) and shifted caregiver behaviour increasing their participation in educational play activities with the child ($d = 0.3$; $p < 0.01$). I also explore the constraints that limit caregivers' behavioural change. I find that treatment effects are not significant among the poorest caregivers and that they exhibit a positive wealth gradient. Further, I find evidence suggesting that caregiver beliefs condition their responses to treatment. **(300 words)**

Acknowledgements

I would like to express my gratitude to the people who helped and supported me throughout the DPhil experience and contributed towards this research effort.

First of all, I would like to thank my supervisors, Professor Stefan Dercon and Dr. Caine Rolleston, for their patient advice and keen insight. I am also grateful to my assessors who provided very useful comments on earlier versions of this work: Dr. Andreas Georgiadis (who assessed my transfer paper), Dr. Imane Chaara (who assessed my transfer and confirmation papers), Professor Douglas Gollin (who assessed my confirmation paper and was part of my final examination committee) and Professor Paul Glewwe (who examined the complete version of the thesis and was also part of my final examination committee).

I am also thankful for fruitful discussions to my good friends and colleagues at the Department of International Development: Carmen Contreras-Romero, Maria del Carmen Franco and Marco Haenssger. Valuable comments were also received from participants at the Queen Elizabeth House Development Economics Seminars, seminars at Universidad del Pacifico's the Research Centre, the 49th Annual Conference of the Canadian Economic Association, the 12th Midwest International Economic Development Conference, and the Latin American and Caribbean Economic Association 2015 Annual Meeting.

I would also like to acknowledge financial support from Universidad del Pacifico, the St Cross College Travel and Research Fund, the Santander Academic Travel Awards, and the Queen Elizabeth House Travel Fund. Special thanks also to the Young Lives Study who shared with me their School Survey, to the *Cuna Mas* Programme who trusted me with their follow-up study on parenting practices, and to Lorenzo Oimas for his outstanding professionalism in organizing fieldwork.

Finally, I would like to dedicate this thesis to my family. To my parents, Juan and Francisca, for their constant encouragement and, especially, to my wife Mapi and my daughter Adriana. Mapi and Adriana shared with me the adventure of studying at Oxford and their company surely made this a much more enriching experience. They are a constant source of inspiration, support and love. A very special thanks to them for bearing with me despite the odd state of mind provoked by the task of putting together this document.

Contents

Chapter 1 <i>Introduction: Inequalities in Child Development in Peru</i>	10
1.1 Developmental outcomes and socioeconomic status in Peru: significant urban/rural gaps emerge early and do not fade away	10
1.2 Theory and evidence on the determinants of cognitive development: differences in early parental influences play a role	18
1.3 Evidence revisited and research questions	27
1.3.1 The evidence revisited	27
1.3.2 Research questions, overview of the methods and expected contributions	32
 Chapter 2 <i>Explaining the Urban-Rural Gap in Cognitive Achievement in Peru: The role of Early Childhood and School Influences</i>	 38
2.1 Introduction and motivation	38
2.2 Educational inequality in Peru	44
2.3 Determinants of cognitive skill and empirical challenges for a gap decomposition	49
2.3.1 The production function of skill	49
2.3.2 Estimating production function parameters	52
2.3.3 Assigning the contribution of input determinants	57
2.4 Decomposition strategy	59
2.4.1 Empirical specifications and assumptions	59
2.4.2 Urban/rural gap decompositions	65
2.4.3 Data sources and variables	70
2.5 Results and discussion	79
2.6 Concluding remarks	89
Appendix 2.1 Consistent estimation of the contribution of an input category	94
Appendix 2.2 Coefficient estimates for the four empirical specifications	96
Appendix 2.3 Estimation and decomposition results after instrumenting lagged PPVT scores with CDA test scores	102
Appendix 2.4 Main decomposition results using the first two principal components of each school quality dimension	105

Chapter 3 <i>Linear Decompositions of Cognitive Achievement Gaps:</i>	
<i>A Cautionary Note and an Illustration Using Peruvian Data</i>	106
3.1 Introduction and motivation	106
3.2 Decomposing achievement gaps in developing countries: arbitrary choices and potential biases.....	110
3.2.1 Choosing a decomposition strategy and interpreting its components	113
3.2.2 Building up categories and assigning contributions	118
3.3 The production function of skill and families' choices regarding its inputs.....	122
3.4 The cognitive skill gap, empirical specifications and decomposition strategies...	130
3.4.1 The risk of bias under the standard decomposition rule	131
3.4.2 An alternative decomposition strategy: Oaxaca-Blinder with a twist	140
3.5 An illustration using Peruvian data	143
3.5.1 Data sources and variables	144
3.5.2 Decomposition results and discussion	146
3.6 Concluding remarks.....	156
Appendix 3.1 From the cumulative to the value added specification with age-dependent parameters.....	166
Appendix 3.2 Coefficient estimates for the eight specifications involved	167
 Chapter 4 <i>Effects of Early Childhood Development Interventions on</i>	
<i>Parental Behaviour: Evidence from a Home-Visiting Programme in Peru</i>	173
4.1 Introduction and motivation	173
4.2 Home-visiting ECD interventions, the home environment and parental behaviour: evidence from the developing world.....	176
4.3 The intervention and research design	182
4.3.1 <i>Cuna Mas</i> home-visiting programme	182
4.3.2 Evaluation design	183
4.3.3 Survey instruments	187
4.3.4 Balance between treatment and control groups at baseline	189
4.4 Treatment effects on the quality of the home environment and parental behaviour	191
4.4.1 Estimation of treatment effects.....	191
4.4.2 Effects on the quality of the home environment	192
4.4.3 Effects on caregiver behaviour.....	194
4.5 Concluding remarks.....	200

Appendix 4.1 Studies that have evaluated interventions with a home-visiting component	203
Appendix 4.2 Districts and regions involved in the study	205
Appendix 4.3 Quality of home environment, household, caregiver and child characteristics in the districts considered for the evaluation	207
Appendix 4.4 Survey questionnaire and characterization of the quality of the home environment	208
Appendix 4.5 Complete regression results for treatment effects on the quality of the home environment	214
 Chapter 5 Constraints Influencing the Effects of Home-Visiting Interventions on Parental Behaviour	215
5.1 Introduction	215
5.2 Heterogeneity of treatment effects: the role of caregivers' wealth and educational attainment	218
5.3 Measurement and treatment effects on caregivers' beliefs regarding the importance of parent-child interactions for child development	222
5.4 A simple time allocation model to explore the mechanisms	227
5.5 Concluding remarks	238
Appendix 5.1 Results related to caregivers' beliefs regarding the importance of parent-child interactions for child development	241
Appendix 5.2 Treatment effects on the variety of play material	247
Appendix 5.3 Treatment effects on the variety of play activities according to the scores obtained in the caregiver beliefs scale	248
 Chapter 6 Conclusions, policy implications, limitations and avenues for future research	249
6.1 Main conclusions and policy implications	249
6.2 Limitations and avenues for future research	255
 References	260

List of Tables

Table 1.1 Peru: standardized PPVT scores ^(a) at 5 and 8 years of age by wealth quartile ^(b) and geographical domain	13
Table 1.2 Peru: proportion of stunted children ^(a) at 1, 5 and 8 years of age by wealth quartile ^(b) and geographical domain	14
Table 1.3 Peru: average monthly per-capita expenditure (2014 Nuevos Soles in Lima)	17
Table 2.1 Children's cognitive achievement, years of schooling, and school and teacher characteristics by student's wealth ^(a) quartile (Q)	48
Table 2.2 Empirical specifications	59
Table 2.3 Variable categories related to each empirical specification and their contribution to the urban-rural gap in cognitive skill	69
Table 2.4 Time-structure and sample sizes of the relevant Young Lives databases	70
Table 2.5 Description of the variables used in the empirical specifications	75
Table 2.6 Descriptive statistics and urban-rural gaps.....	77
Table 2.7 Normalized contributions to the urban-rural gap in cognitive skill at age 8 (% of urban-rural gap).....	80
Table 3.2 Four different two-fold Oaxaca-Blinder decompositions	115
Table 3.3 Categories and variables included under the standard decomposition rule.....	137
Table 3.4 Categories and variables included under the alternative decomposition strategy	141
Table 3.7 Normalized contributions to the urban-rural gap in cognitive skill at age 8 (% of urban-rural gap).....	148
Table 3.1 Studies that have attempted a linear decomposition of the observed cognitive achievement gap between two groups of children	159
Table 4.1 Outcome variables and household, caregiver and child characteristics at baseline	190
Table 4.2 Treatment effects on the quality of the home environment.....	193
Table 4.3 Treatment effects on caregivers' interaction with child and on the variety of play activities offered by the caregiver	196
Table 4.4 Effects on the variety of play activities offered to the child (number of types of play activities)	198
Table 4.5 Treatment effects on each type of play activity (proportion of caregivers)	200
Table 5.1 Treatment effects by household wealth and caregiver education quartiles on the variety of play activities offered by the caregiver at least twice per week.....	220
Table 5.2 Interactions between treatment status, household wealth and caregivers' educational attainment	220
Table 5.3 Treatment effects on caregivers' beliefs regarding the importance of parent-child interactions for child development	226
Table 5.4 Treatment effects by household wealth and caregiver education on caregivers' beliefs regarding the importance of parent-child interactions for child development (average score)	226

List of Figures

Figure 1.1 Peru: standardized scores for learning outcomes in 2nd grade (age 8) ^(a)	15
Figure 1.2 Peru: proportion of stunted 0-5 year-old children (chronic malnutrition) ^(a)	16
Figure 1.3 Peru: caregiver characteristics of urban and rural children at 5 years of age (point estimates and 95% confidence intervals) ^(a)	24
Figure 1.4 Peru: Variety and frequency of play activities and variety of play material offered to children in rural households (gradients by household wealth and caregiver educational attainment)	25
Figure 1.5 Peru: Projected urban/rural differences in the variety and frequency of play activities and the variety of play material offered to children (point estimates and 95% confidence intervals) ^(a)	27
Figure 1.6 Peru: Differences between urban and rural public schools ^{(a) (b)}	29
Figure 2.3 Normalized contributions to the urban-rural gap in cognitive skill at age 8 (point estimates and 95% confidence intervals)	81
Figure 2.4 Normalized contributions to the urban-rural gap in cognitive skill at age 8 - excluding school inputs from the school survey	87
Figure 2.5 Normalized contributions to the urban-rural gap in cognitive skill at age 8 - excluding school inputs from the school survey and working with the complete sample	88
Figure 2.1 Proportion of second grade students with a satisfactory attainment in reading comprehension (urban and rural point estimates and 95% confidence intervals)	92
Figure 2.2 Urban-rural differences in cognitive achievement (standardized Peabody Picture Vocabulary Test scores)	93
Figure 3.1 Normalized contributions to the urban/rural gap in cognitive skill at age 8 (point estimates and 95% confidence intervals)	150
Figure 5.1 Treatment effects by household wealth and caregiver educational attainment	221
Figure 5.2 Relationship between caregivers' beliefs regarding the importance of parent-child interactions for child development (average score) and caregivers' household wealth and education	227
Figure 5.3 Two scenarios for treatment effects on childcare time	238

Chapter 1

Introduction: Inequalities in Child Development in Peru

1.1 Developmental outcomes and socioeconomic status in Peru: significant urban/rural gaps emerge early and do not fade away

Differences in developmental outcomes between children of dissimilar socioeconomic backgrounds are particularly significant in the developing world (Grantham-McGregor et al., 2007; Walker et al., 2007). These differences or gaps emerge early in the life of children and tend to persist throughout their school years (Heckman, 2006, 2007; Paxson and Schady, 2007a; Schady et al., 2014). These gaps then translate into inequalities in terms of labour market outcomes and other indicators related to their personal wellbeing. Numerous studies have found a strong causal relationship between developmental indicators during childhood and later-life outcomes such as post-secondary schooling, employment status, wages, and participation in crime (Almond and Currie, 2011; Cunha et al., 2006). Therefore, persistent setbacks in developmental outcomes among children from disadvantaged groups explain one of the channels through which poverty is inter-generationally transmitted and inequality is sustained.

Peru is no exception to the presence of these early forms of inequality. The literature distinguishes three basic domains of child development: sensorimotor, cognitive and socio-emotional (Boyden and Dercon, 2012; Grantham-McGregor et al., 2007). Peru exhibits positive and significant socioeconomic gradients in several indicators related to these domains. In tables 1.1-1.2 and figures 1.1-1.2, I present evidence on differences by

household wealth and geographical domain in two indicators related to the cognitive domain and in a widely acknowledged risk factor¹ for child development.

Tables 1.1 and 1.2 present longitudinal evidence for a large sample of Peruvian children at ages 1, 5 and 8.² This type of information is particularly useful to determine if inequalities in child development emerge early in life and if these gaps tend to grow or diminish as children age. As documented in other recent studies,³ Table 1.1 confirms that gaps in a well know measure of cognitive development emerge before children enter school and exhibit little variation as they grow older and progress into primary education. Differences in test scores between the first and fourth quartile of the wealth distribution in the Young Lives sample are around 1.5 standard deviations at both 5 and 8 years of age. Interestingly, Table 1 also reveals that the urban/rural gap is very significant (amounting to, at least, 1 standard deviation) and remains practically unchanged between ages 5 and 8.

Table 1.2 documents the proportion of children with growth retardation (that is, with a height for age ratio below the norm). This is a widely used indicator of chronic malnutrition or stunting. Stunting, in turn, is a good indicator of poor child development and a widely acknowledged risk factor (Grantham-McGregor et al., 2007). Results refer to the same sample of children considered for the cognitive skill measure presented in Table 1.1 at 1, 5 and 8 years of age. By age 1, the difference in the proportion of stunted children between the urban and rural domains amounted to nearly 25 percentage points. This figure is similar

¹ Risk factors refer to biological and psychosocial hazards that can jeopardize child development (Walker et al., 2007).

² This is based on the information collected for the younger cohort of the Young Lives Study in Peru. Young Lives is an international study of childhood poverty, following 12,000 children in 4 countries (Ethiopia, India, Peru and Vietnam) over 15 years. The first round was collected in year 2002. In Peru, it follows two cohorts of children: the younger cohort (aged 1 in 2002) and the older cohort (aged 8 in 2002). The Peruvian sample has a slight pro-poor bias as it excludes households belonging to the richest 5% of districts. Therefore, socioeconomic gradients must be even steeper than those reported in tables 1.1 and 1.2. Because excluded rich districts are only urban, the sample can be regarded as representative of the rural domain.

³ Recent work by Schady et al. (2014) has documented socioeconomic gradients for cognitive development in five Latin American countries including Peru. They found steep wealth gradients in a common measure of cognitive development in all five countries, and confirmed that these gaps are well established by age three and remain practically invariant once children enter school.

to the gap between the first and the third quartile of the Young Lives wealth distribution, and shows no evidence of subsiding as children grow older.

Figures 1.1 and 1.2 complement the abovementioned evidence with urban/rural gaps estimated using nationally representative surveys. Figure 1.1 presents the recent evolution of learning outcomes attained by second grade students (around age 8). Figure 1.2 presents chronic malnutrition figures for 0-5 year-old children. This information is useful in two ways. First, it confirms the existence of significant urban/rural differences in the cognitive development of Peruvian children as well as in a risk factor for child development in general. Second, it reveals another form of persistence in early childhood inequality. Tables 1.1 and 1.2 have shown that socioeconomic differences in child development emerge early and do not subside as children grow older. Figures 1.1 and 1.2 reveal that inequalities present in particularly relevant age groups have not subsided in recent years despite the significant economic growth rates attained by Peru.⁴ These gaps are still significant today and should draw the attention of policymakers.

⁴ According to the National Institute of Statistics, in the period between 2005 and 2013, per capita GDP grew at an average rate of 5.4% per year.

Table 1.1

Peru: standardized PPVT scores^(a) at 5 and 8 years of age by wealth quartile^(b) and geographical domain

	No. of obs.	Mean		Difference ^(c)	
		5 year-olds	8 year-olds	5 year-olds	8 year-olds
A: Household wealth					
Q1	430	-0.78	0.95		
Q2	433	-0.34	1.53	0.44***	0.58***
Q3	433	0.33	2.09	1.11***	1.14***
Q4	421	0.81	2.34	1.6***	1.39***
B: Geographical domain					
Urban	1,234	0.30	2.02		
Rural	483	-0.76	0.97	1.05***	1.05***
All	1,717	0.00	1.73		

Source: Second and third round of data from the younger cohort of the Young Lives Study. The information was collected in years 2006 and 2009, respectively.

(a) PPVT scores refer to the result obtained in the Peabody Picture Vocabulary Test. This is a test for receptive vocabulary that has a Spanish version adapted for Latin America (Dunn et al., 1986). It is a widely used measure of cognitive achievement as it has a strong positive correlation with several measures of intelligence (Cueto and Leon, 2012). Raw scores were standardized using the mean and standard deviation obtained in round 2.

(b) Based on a wealth index comprising dwelling characteristics, access to basic services and durable goods consumption (Escobal et al., 2003).

(c) For household wealth, differences were computed with respect to the first quartile. For the geographical domain, the difference is urban minus rural.

Significance levels: *** p<0.01, ** p<0.05, * p<0.1

Table 1.2

Peru: proportion of stunted children^(a) at 1, 5 and 8 years of age by wealth quartile^(b) and geographical domain

	No. of obs.	Mean			Difference ^c		
		1 year-olds	5 year-olds	8 year-olds	1 year-olds	5 year-olds	8 year-olds
A: Household wealth							
Q1	430	0.45	0.59	0.37			
Q2	433	0.33	0.42	0.26	-0.12***	-0.17***	-0.11***
Q3	433	0.17	0.17	0.08	-0.28***	-0.42***	-0.29***
Q4	421	0.13	0.10	0.07	-0.33***	-0.49***	-0.3***
B: Geographical domain							
Urban	1,234	0.20	0.22	0.13			
Rural	483	0.44	0.57	0.37	-0.23***	-0.35***	-0.24***
All	1,717	0.27	0.32	0.20			

Source: First, second and third round of data from the Younger Cohort of the Young Lives Study. The information was collected in years 2002, 2006 and 2009, respectively.


(a) A child is considered stunted if she exhibits a height for age z score below -2 using the WHO standard.

(b) Based on a wealth index comprising dwelling characteristics, access to basic services and durable goods consumption (Escobal et al., 2003).

(c) For household wealth, differences were computed with respect to the first quartile. For the geographical domain, the difference is urban minus rural.

Significance levels: *** p<0.01, ** p<0.05, * p<0.1


Figure 1.1
Peru: standardized scores for learning outcomes in 2nd grade (age 8) ^(a)



Source: Evaluacion Censal de Estudiantes 2010-2014, Ministry of Education, Peru.

(a) Reading comprehension and mathematics tests are taken every year by all second grade students in schools with five or more potential test takers. Scores were standardized using the mean and standard deviation obtained in year 2010.

Figure 1.2
Peru: proportion of stunted 0-5 year-old children (chronic malnutrition)^(a)



Source: Encuesta Demografica y de Salud Familiar (ENDES), National Institute of Statistics, Peru.
(a) A child is considered stunted if she exhibits a height for age z score below -2 using the WHO standard.

It is clear from the evidence reported above that, in Peru, differences in household wealth and other socioeconomic status indicators are strongly related to the geographical domain of the household. Nearly a quarter of the Peruvian population (24%) lives in rural areas. However, rural areas concentrate 58% of people belonging to the first quintile of the expenditure distribution and only 4% of those belonging to the upper quintile. Accordingly, average per-capita expenditure among urban households is more than twice that of rural households, and shifting from the rural to the urban domain resembles shifting from some point between the first and second quintile to some point between the fourth and the fifth quintile in the national per-capita expenditure distribution (see Table 1.3).

Table 1.3**Peru: average monthly per-capita expenditure (2014 Nuevos Soles in Lima)**

	1st quintile	2nd quintile	3rd quintile	4th quintile	5th quintile	All
Urban	404	466	541	669	1,087	732
Rural	250	311	410	524	809	358
Peru	303	403	512	651	1,075	642

Source: Encuesta Nacional de Hogares (ENAHOG) 2014, National Institute of Statistics, Peru.

Lower expenditure levels and higher poverty rates observed in the rural domain (46% monetary poverty incidence vs. 15.3% in urban areas (INEI, 2015)) are caused by an economy based on low productivity or subsistence farming. According to the 2014 National Household Survey (2014 ENAHOG), nearly 75% of household heads in the rural domain work in agriculture or livestock farming and 84% of them are self-employed or unpaid family workers.

Urban/rural differences are also striking in terms of dwelling characteristics and access to basic services. According to the 2014 ENAHOG, only 17% of houses in the rural domain have their floor made of finished material, a sturdy roof and brick or plastered walls. The same indicator in the urban area reaches 74%. Access to piped water (64%) and flush toilets or pit latrines (60%) is also limited in the rural domain, whereas in urban areas there is almost universal access (these proportions are as high as 93% and 92%, respectively). All this evidence is consistent with the similarities found between the wealth and urban/rural gradients in child development documented in tables 1.1 and 1.2.

A pertinent question with direct policy implications is: what explains or drives these urban/rural differences in child development? As documented above, the division of Peruvian households into urban and rural residents constitutes a major axis of inequality in terms of both outcomes and access to basic services. Understanding how urban/rural differences in the provision of public services and programmes translate into gaps in

outcomes related to the wellbeing of their inhabitants should provide important policy lessons. This relates to the broader questions of which are the determinants of the developmental outcomes discussed above, and how effective are policy interventions in changing these determinants.

In the next section I summarize theoretical knowledge and empirical evidence regarding the determinants of child development. The emphasis will fall on early parental influences, for which I also document urban/rural differences using Peruvian data. I then use this discussion to motivate the specific research questions proposed for the thesis. The analysis will focus on the domain of cognitive development. This domain plays a prime role in the process of human capital accumulation and has strong causal links with labour market and other later life outcomes (Cunha et al. (2006); see Castro et al. (2012) and Diaz et al. (2013) for the Peruvian evidence).

1.2 Theory and evidence on the determinants of cognitive development: differences in early parental influences play a role

Several disciplines have contributed to the understanding of cognitive development in children. The most influential is Developmental Psychology and, more recently, Cognitive Neuroscience. Economics has also made theoretical and empirical contributions through the literature on the “technology of skill formation”.

Two significant contributions made by Developmental Psychology are the notions that the acquisition of cognitive skill is a cumulative process and that the environment matters. The first notion can be related to the stage theories of developmental psychology, among which Jean Piaget’s theory of cognitive development stands as the most well-known and influential (Flavell, 1982).

The notion that environments or contexts matter stems from the contributions first introduced in the 1970s by Urie Bronfenbrenner. These have been structured around what the author has more recently called the “bioecological model of human development” (Urie Bronfenbrenner and Morris, 2006). According to this model, individuals develop in interaction with their immediate environment. These interactions are referred to as “proximal processes”. The form, content and effect of these processes are, in turn, determined by the characteristics of the person, the environment (or context) and of the time period in which they take place.

Recent advances made by Cognitive Neuroscience have allowed a better understanding of why particular elements of the environment matter for cognitive development. This is related to the identification of mechanisms linking experience to behaviour through molecular changes in gene expression and brain circuits. Studies reviewed in Farah and Hackman (2012) have been able to: (i) identify which specific subsystems involved in cognitive processes are affected by harsh socioeconomic conditions (i.e. the language and executive function subsystems⁵); and (ii) isolate particular elements of these environments (e.g. poor cognitive stimulation and stress) and elucidate the chemical pathways through which they affect the structure and functioning of the brain.

Piaget’s idea of orderly and progressive universal stages of maturation in human development has not remained unchallenged. For example, Flavell (1982) presented evidence of heterogeneous maturity levels, complexity and styles of thinking in the mind of a child at a particular age. This heterogeneity is not consistent with a stage-like development of human cognition. Burman (1994), challenged the “normalisation” and “naturalization” of

⁵ The development of these two subsystems consistently exhibits a significant socioeconomic gradient. The executive function system “enables flexible responding in situations where the appropriate response may not be the most routine or attractive one, or where it requires maintenance or updating of information concerning recent events” (Farah and Hackman, 2012; p. 310).

human development. She criticized the establishment of norms that reflect an “adequate” development, and examined the process by which these ideas were given scientific legitimation by paralleling the study of the mind to the study of the body.

Several authors have also argued against over-deterministic views of how childhood experiences affect later-life outcomes. For example, Yaqub (2002) reviewed evidence related to human physiological development which showed that negative influences do not always cause developmental setbacks (indicating heterogeneous levels of resilience) and, even if they do, that these setbacks can be reversed (indicating plasticity in developmental pathways).

In sum, the possibility of talking about ordered and universal stages in cognitive development is still open to debate, and there is no conclusive evidence of a critical period for human development. There is, however, strong empirical evidence to support the importance of early influences for outcomes observed later in life, and the existence of sensitive periods when some skills are more malleable (see Almond and Currie (2011) and Cunha and Heckman (2007) for reviews). This evidence provides the foundation for relating skill measures taken at a particular age not only to contemporaneous influences but also to the skill measures and environments observed earlier in the life of the child, and for allowing these environments to exert varying degrees of influence depending on the age transition when they occurred.

These notions are closely related to the contributions made by Economics through the concept of “skill formation technology”. As suggested by its name, the “technology” or “production function” approach to the analysis of skill formation understands the determinants of skill as inputs and seeks to represent the manner in which these are combined to produce a particular skill. Within the economic literature, efforts to empirically

assess the importance of diverse influences for the acquisition of cognitive skills can be traced back to the seminal work of Leibowitz (1974). More recent contributions can be found in the works of Todd and Wolpin (2003, 2007) and Cunha and Heckman (2007, 2008).

On the theoretical front, the economic literature has introduced concepts appropriate to characterize the production of an output with a cumulative nature such as cognitive skill. It has also contributed models explaining how different socioeconomic conditions translate into different quantities of skill inputs through the decision making process of families. In particular, in Cunha and Heckman (2007), the authors develop a multi-period model for childhood development and introduce three sets of concepts to explain the main features of the skill formation technology: (i) critical and sensitive periods; (ii) self and cross productivity; and (iii) dynamic complementarity.⁶

On the empirical front, this literature has placed special emphasis on: (i) identifying the assumptions required for consistent estimation of the effects of particular inputs or influences on skill by acknowledging their endogenous nature; (ii) testing the relationship between cognitive and socio-emotional skills and; (iii) testing the existence of sensitive periods for parental investments (Cunha and Heckman, 2007, 2008; Cunha, Heckman, and Schennach, 2010; Todd and Wolpin, 2003, 2007). Findings based on the US National Longitudinal Survey of the Young (NLSY/79) show that skill formation is a cumulative process, that socio-emotional skills affect cognitive skill, and that cognitive skill is

⁶ A critical period for a particular skill is the only stage in which influences or inputs can have an effect. A sensitive period is one in which inputs are especially effective or productive. Self-productivity is a property by which skills attained at one stage augment the skills attained at later stages. From the point of view of a particular skill, this means that the ability attained at one stage persists in future stages (or that skill does not fully depreciate from one period to the other). From the point of view of multiple skills, this means that a particular skill attained at a particular stage can also foster the acquisition of other skills in subsequent stages (cross-productivity). Dynamic complementarity refers to the property by which skill produced at one stage of the life cycle raises the effect of influences that occur in later stages on the same skill (i.e. they increase the productivity of later inputs).

particularly sensitive to parental investments during early childhood. Efforts to replicate this in the developing world have confirmed the importance of early parental investments and the fact that cognitive and socio-emotional skills are related (Helmers and Patnam, 2011; Lopez-Boo, 2009). Studies have also found strong causal links between time spent in educational activities with parents and cognitive skill formation among children (Fiorini and Keane, 2014).

The idea that parenting factors matter is not new, nor exclusive to Economics. It has been extensively discussed how influences exerted by parents have a prime role for early child development and how parenting factors such as sensitivity and cognitive stimulation potentially mediate links between socioeconomic status and child outcomes (see Huberman and Mendelsohn (2012) and Pungello and Gardner-Neblett (2012) for recent reviews). To this discussion, the skill formation literature from Economics has contributed more evidence on the causal link between parental actions during early childhood and children's cognitive development, as well as structure explaining these actions as the result of the rational behaviour of parents concerned with the future wellbeing of their children (hence the term "parental investments").

Based on the notions and evidence documented above, it is reasonable to trace back socioeconomic gradients in children's cognitive development to differences in influences exerted and provided by caregivers of dissimilar socioeconomic backgrounds. It is already well established that several aspects of parenting,⁷ including parenting practices, vary with the socioeconomic status of parents (Hoff et al., 2002). Moreover, the lack of early cognitive stimulation has been acknowledged as one of the four key risk factors for child

⁷ The literature distinguishes at least three aspects of parenting: (i) goals; (ii) styles; and (iii) practices. Goals refer to the objectives towards which parents orient their efforts. Styles comprise parental attitudes towards children. Parenting practices include the behaviours parents produce in interactions with their children (Hoff et al., 2002).


development in developing countries⁸ (Walker et al., 2007). These authors reviewed 16 intervention studies that assessed the effect of cognitive stimulation on children in developing countries. They found significant effects in all but one intervention. Importantly, they also found that only 10-41% of parents provide cognitively stimulating materials to their child and only 11-33% of parents actively involve their children in cognitively stimulating activities.

Consistent with this evidence and with the urban/rural gaps in cognitive development documented in the previous section, caregivers living in the urban and rural domains in Peru differ in various aspects that can influence the quantity and quality of cognitive stimulation that they provide to their children. Figure 1.3 documents urban/rural differences in several caregiver characteristics that, according to the literature, are positive predictors of parenting practices: education, wealth and aspirations. It uses the same sample of children that yielded the urban/rural gaps in cognitive skill presented in Table 1.1. It is clear from Figure 1.3 that urban caregivers are more educated, wealthier, and are more likely to desire that their children attend university higher education.

⁸ The other four are stunting, iodine deficiency and iron deficiency anaemia.

Figure 1.3

Peru: caregiver characteristics of urban and rural children at 5 years of age (point estimates and 95% confidence intervals)^(a)



Source: Second round of data from the Younger Cohort of the Young Lives Study. This information was collected in year 2006. The number of observations is 1,717 (1,234 urban and 483 rural households).

(a) Confidence intervals were built using robust standard errors.


(b) The household wealth index combines information on dwelling characteristics, access to basic services and availability of durable goods (Escobal et al., 2003).

Results from a large survey on childrearing practices conducted in the rural domain⁹ confirm that the socioeconomic characteristics considered above are significant predictors of parenting practices. This survey allows one to build three indicators related to the level of cognitive stimulation provided to the child: the variety and frequency of play activities offered to the child and the variety of play material to which the child has access. Figure 1.4 confirms the existence of positive gradients by caregiver education and wealth in all three indicators.

⁹ This survey is the baseline study for the rural component of the early childhood development programme *Cuna Mas*. The sample was collected in year 2013 and comprises information on 5,355 caregivers living in 360 rural communities across 12 regions.

Figure 1.4

Peru: Variety and frequency of play activities and variety of play material offered to children in rural households (gradients by household wealth and caregiver educational attainment)



Source: Baseline study for the rural component of programme *Cuna Mas*. It comprises information on 5,355 caregivers living in 360 rural communities across 12 regions. It was collected in year 2013.

(a) Number of types of play activities offered to the child at least once per week. These include: (i) reading books; (ii) telling stories; (iii) singing songs; (iv) taking child outside the house; (v) playing with child and her toys; (vi) drawing or painting; and (vii) naming objects or numbers.

(b) Refers to the average number of days per week in which play activities are carried out.


(c) Number of types of play material present at home. These include: (i) children books; (ii) material to draw or paint; (iii) cuddly or role playing toys; (iv) push or pull toys; (v) building blocks or puzzles; and (vi) toys to play music.

(d) The household wealth index combines information on dwelling characteristics, access to basic services and availability of durable goods (Escobal et al., 2003).

Information on parenting practices from urban caregivers is not available. However, it is possible to use the predictors considered in Figure 1.4 to extrapolate parental behaviour in urban households. In particular, I projected the variety and frequency of play activities and the variety of play material in the urban and rural domains using OLS coefficient estimates obtained from the rural survey and caregiver characteristics from the same sample of children that produced the cognitive skill gaps presented in Table 1.1. Results are presented in Figure 1.5. They illustrate how urban children in Peru exhibit a positive and significant advantage in terms of their involvement in cognitively stimulating activities and their access to educational play material.

Figure 1.5

Peru: Projected urban/rural differences in the variety and frequency of play activities and the variety of play material offered to children (point estimates and 95% confidence intervals)^(a)



Sources: Baseline study for the rural component of programme *Cuna Mas* and second round of data from the Younger Cohort of the Young Lives Study.

(a) Projection uses coefficient estimates obtained from the *Cuna Mas* baseline study (5,355 rural caregivers) and covariate values from the second round of data of the Younger Cohort of the Young Lives Study (1,717 children; 1,234 urban, 483 rural). Predictors include: caregiver education, household wealth, household size, caregiver age, caregiver mother tongue, if caregiver works outside the household, and if caregiver is the child's mother. Confidence intervals were built using robust standard errors.

1.3 Evidence revisited and research questions

1.3.1 The evidence revisited

The Peruvian evidence surveyed above presents us with urban/rural gaps in cognitive development that emerge at an early age and persist as children grow older. I have also reviewed the theoretical reasons and empirical evidence that support the claim that these

gaps are, at least in part, caused by differences in the quantity and quality of parenting factors present during early childhood.

Two obvious questions that follow and have clear policy implications are: (i) if there are other influences, besides those exerted by parents during early childhood, that determine the emergence of these gaps in child development; and (ii) how amenable are early childhood parenting factors to policy intervention. As will be discussed below, the literature still does not offer a satisfactory answer to these questions and they therefore represent the main motivation for this thesis.

Regarding the first question, the school environment is a good candidate to look for other influences potentially relevant for child development. In fact, there is a vast literature on school effectiveness presenting numerous examples from the developing world of how school and teacher characteristics can make a difference in learning outcomes that are closely related to the domain of cognitive development (see Glewwe et al. (2013) for a recent review).


In terms of education indicators, Peru has achieved high and fairly homogenous rates of education coverage but still exhibits poor and highly unequal indicators in terms of school and teacher characteristics (Beltran and Seinfeld, 2012; Cueto, Guerrero, et al., 2014). The net enrolment rate in primary education of 86.2% is now close to 90% and exhibits no significant difference between urban (87.8%) and rural (82.6%) areas.¹⁰ Urban and rural domains, however, differ significantly in terms of teacher and school characteristics potentially relevant for learning. A recent survey of public schools¹¹ confirms that urban

¹⁰ These figures were calculated using the results of the 2014 National Household Survey (*Encuesta Nacional de Hogares* - ENAHO).

¹¹ According to the 2014 enrolment records of the Ministry of Education, public schools serve nearly three-quarters of basic education students. Private schools operate mainly in urban areas, where they concentrate 31% of enrolment in basic education. The supply of educational services in rural areas is almost entirely public, serving 97% of the students.

institutions have more qualified and experienced primary education teachers, have more opportune access to learning materials, and offer classrooms in better physical conditions (see Figure 1.6).

Figure 1.6
Peru: Differences between urban and rural public schools ^{(a) (b)}



Source: Encuesta Nacional de Instituciones Educativas (ENEDU, 2014). National Institute of Statistics, Peru.
 (a) Classes start in March. Timely Access implies that learning materials were received before May.
 (b) Adequate infrastructure implies: (i) roof has no holes or fissures; (ii) windows have complete and undamaged glasses; and (iii) the classroom has electric lights.

It is worth noticing that the existence of significant differences between urban and rural schools does not, by itself, imply that these play a role in shaping the developmental gaps documented before. In fact, that the prevalence of cognitive achievement gaps that emerge during early childhood remains practically unchanged once children start school apparently suggests that school influences play only a subsidiary role in the emergence of these gaps. This pattern is present in the Peruvian data and has been documented by other authors for

several countries (both developed and developing) to highlight the importance of concentrating investment efforts during early childhood (Heckman, 2006; Schady et al., 2014).

In addition, studies that have addressed the role of schools in the skill formation technology have found inconclusive results using information from both developed (Todd and Wolpin, 2007) and developing countries (Helmert and Patnam, 2011).¹² Other studies that have directly attempted a comparison between the importance of school and family influences for learning outcome gaps in the developing world have found results that lack robustness. (Arteaga and Glewwe, 2014; Hernandez-Zavala et al., 2006; McEwan, 2004; McEwan and Trowbridge, 2007). For example, estimates of the contribution of schools to the difference in learning outcomes between indigenous and non-indigenous children in Guatemala vary between 17% (Hernandez-Zavala et al., 2006) and 70% (McEwan and Trowbridge, 2007).

Regarding the question on the extent to which parenting factors are amenable to policy intervention, the literature on early childhood development offers strong evidence showing that interventions that focus on parent-child interactions can produce positive results in child outcomes (Attanasio et al., 2013; Baker-Henningham and Lopez-Boo, 2010; Engle et al., 2007; Engle et al., 2011; Walker, 2011). There is, however, no unequivocal indication of how these interventions work. These interventions seek to improve parenting practices by demonstrating play activities and child-care routines to caregivers over a period of 1 to 2 years. Demonstrations are usually carried out through home visits and in some cases they also involve the provision of materials such as books and toys. These activities typically

¹² In Todd and Wolpin (2007), the authors did not find evidence of a significant relation between school input measures and cognitive skill test scores using the NLSY. They acknowledge that school input coefficients were not precisely estimated and that the data on school quality was probably too noisy. In Helmert and Patnam (2011), the authors used data from India and considered the type of school management as a potential predictor of cognitive skill. They did not find robust results regarding the advantage of studying in a private school. It is worth noticing that these studies relied on datasets with rich longitudinal information at the child and household level but that lack the same degree of detail at the school level.

involve the child, so measured outcomes include the effect of direct stimulation which cannot be disentangled from the effect that is potentially mediated through a change in parental behaviour.

In this respect, the abovementioned literature presents scant evidence on the effect of these interventions on parenting practices related to the provision of cognitive stimulation¹³ and has paid little attention to the mechanisms through which these programmes can affect parental behaviour. The literature places strong emphasis on the transfer of parenting skills to caregivers (“parent education”) highlighting the importance of curricula, training opportunities for home visitors and parents, and the use of active strategies to ensure caregiver participation (e.g. role play or coaching) (Engle et al., 2011; Huberman and Mendelsohn, 2012). This carries the implicit assumption that parents’ lack of knowledge of and ability to engage in particular activities with their children is the only relevant constraint. This assumption is probably too strong as parenting practices are determined by diverse influences and families focalized for these types of interventions face constraints in many dimensions.

Following Hoff et al. (2002), evidence suggests that differences in parenting between caregivers of dissimilar socioeconomic backgrounds can be traced back to differences in their circumstances (or external factors) and also to intrinsic or internal dissimilarities. The above mentioned literature focuses on very few aspects of these potential sources of heterogeneity. In fact, enhancing skills so that parents can engage in more educational play activities with their children can be understood as a shift in their internal characteristics, in particular, in their knowledge of and ability to perform certain activities. In addition, the

¹³ In Attanasio et al. (2013), Baker-Henningham and Lopez-Boo (2010), Engle et al. (2007), Engle et al. (2011) and Walker (2011), the authors review the effects of 17 home-visiting interventions on child outcomes. All of them report positive results in at least one domain of child development. However, only 9 studies measured changes in the quality of the home environment and just 3 reported positive effects in the stimulation opportunities offered to the child by the caregiver.

provision of books and toys can be characterized as a change in their material circumstances.

Parents' beliefs and goals also fall within the intrinsic characteristics that determine their parenting practices. These aspects have received much less attention in the literature on early childhood development as potential mediators or moderators of interventions focused on parent-child interactions.¹⁴ This is despite evidence showing that parents' beliefs regarding their role in their children's development varies with socioeconomic status and appears to partially explain the relation between socioeconomic status and parenting practices (see Hoff et al. (2002) for a review).

1.3.2 Research questions, overview of the methods and expected contributions

In the previous section I proposed two questions based on the theory and empirical evidence reviewed. In particular: (i) if there are influences besides those exerted during early childhood that determine the emergence of gaps in child development; and (ii) how amenable are early childhood parenting factors to policy intervention.

If we revisit these questions after the preceding discussion, it is clear that: (i) there is still an unresolved problem regarding the relative importance of schools and early childhood influences when explaining the occurrence of unequal developmental outcomes between children of different socioeconomic backgrounds; and (ii) there is limited evidence on the effect of policy interventions on parenting practices related to the provision of cognitive stimulation and little attention has been paid to the mechanisms behind caregivers' behavioural change.

¹⁴ In a recent study of a parenting intervention in Colombia, Attanasio et al. (2015) acknowledge that such interventions can alter parental investment decisions by making them aware of the importance of early investments. This and other potential channels for behavioural change such as knowledge of specific practices, however, were not explicitly considered in their analysis.

Motivated by this review and by the prominence of developmental gaps between urban and rural children in Peru, this thesis seeks to contribute to the literature on skill formation and child development by addressing the following specific research questions:

1. What is the importance of school influences vis-à-vis that of the early learning and care environment when explaining inequalities in cognitive development between urban and rural children in Peru?
2. Can an intervention focused on parent-child interactions working at scale in rural Peru deliver an improvement in the quality of the home environment?
3. Can an intervention focused on parent-child interactions working at scale in rural Peru change parental behaviour so as to increase the amount of stimulation offered to the child?
4. What constraints faced by caregivers limit the effect of this intervention on caregiver engagement in educational play activities with the child?

By proposing to answer these four research questions, this analysis seeks to make contributions not only in terms of new evidence but also in terms of the strategies and methods employed to produce this evidence, especially those pertaining to the first research question. As discussed below, these strategies and methods are closely related to the role that Economics has played when explaining the process of cognitive skill formation and parenting practices.

On the theoretical front, Economics has provided a framework to organize the decision making process of families pertaining to child development. The key ingredients of this framework are parental preferences, a budget constraint and a production constraint (Foster, 2002). Preferences capture the value that parents place on child outcomes and other goods.

The budget constraint relates family resources with the prices and quantities of goods and services consumed. These can be directly valued by the family or used as inputs to produce child outcomes and other goods. Finally, the production constraint relates key inputs with child outcomes. A good example of this constraint is the technology of cognitive skill discussed above.

In Foster (2002), the author discusses four features of this framework. In particular, the economic framework: (i) highlights the identification problems related to the effects of particular influences on child development because it explicitly recognizes that many of these influences are part of the same choice process (what Foster calls the “problems of self-selection”); (ii) it recognizes that families respond to changes in resources and incentives that can be affected by policy action; (iii) it can explicitly allow for interactions or complementarities in the production of child outcomes; and (iv) it provides an expanded understanding of the context in which the child develops and a richer view of the effect of family resources.

The first feature of the economic model reminds us to exercise caution when interpreting partial correlations as causal effects. Because it makes explicit the process that determines child development inputs, an economic model may also contribute a way to circumvent the identification problem in hand by, for example, suggesting a potential instrumental variable.

This explicit consideration of the choice process that delivers the inputs of child development is also related to the fourth feature mentioned above. In fact, an economic model of family production places less emphasis on the specific interactions that produce child development (Bronfenbrenner’s “proximal processes”) and concentrates on characterizing the environment in which this development occurs. This environment can be divided into those elements with which the child interacts (the inputs of child development)

and those elements which determine the quantity and quality of the inputs. These latter elements of the environment typically comprise families' preferences, the prices they face and the resources available to them. In the solution to an economic model, the inputs of child development end up being a function of these preferences, prices and resources.

In principle, variables reflecting inputs (such as play activities) and input determinants (such as family income) are all potentially good predictors of child outcomes. An economic model of family production proposes a convenient way of understanding their relation and organizing their role in the production of child outcomes. This insight is extremely useful when interpreting partial correlations between child outcomes and family characteristics and when choosing controls to try to circumvent an omitted variable bias.

The second feature mentioned above relates to the fact that an economic model places policy interventions in the context of family resource allocation and this insight can help assess their effectiveness. Interventions typically provide some valued good directly to their beneficiaries or provide incentives that should lead to some desirable shift in their allocation of resources. Accounting for behavioural change is important in both cases. If policy is directly providing a good or mitigating a constraint, this accounting can help anticipate unintended effects. In the second case, having some idea of the potential mechanisms behind behavioural change is even more important as the intervention is conceived to work through a change in resource allocation decisions.

The merit of the strategy employed to answer the first research question of this thesis stems from the fourth feature of the economic framework discussed above. In particular, the first research question will be addressed by means of a decomposition strategy that explicitly acknowledges how child development inputs (including those provided at school) are determined as a result of family preferences, resources, available goods and services, and

their prices. This framework will prove important to develop an empirical strategy less prone to biases than those employed thus far to compare the importance of school and family influences. This will also serve to shed light on why the results reported in the literature lack robustness.

The second feature of the economic model outlined above will play an important role when addressing the fourth research question and, in particular, when rationalizing why it is that certain caregiver and household characteristics can limit the effectiveness of a parenting intervention. Programmes focused on parent-child interactions seek to induce a change in caregivers' behaviour. This behaviour can be understood in terms of a resource allocation problem. In fact, parental time is a prominent example of a scarce resource that is required for the production of activities involving parents and their children. Therefore, an economic model of family production can provide insights into the mechanisms behind caregivers' behavioural change that have implications for policy design. In fact, understanding the parenting skills provided by the intervention in the context of a resource allocation problem can serve to identify complementary incentives or information that can also be offered to caregivers to enhance the effect of the programme.

Finally, the empirical strategy employed to address the second and third research questions can also be motivated from the economic framework discussed above. As already mentioned, one of the merits of the economic model is that it highlights the selection problem that complicates the identification of the effects of particular influences on child development. This identification problem extends to the estimation of the effects of parenting interventions on caregivers' behaviour. Consistent with this, research questions 2 and 3 will be addressed by means of an experimental design. This should mitigate concerns

regarding potential biases as random assignment of the treatment prevents correlation between treatment status and unobservable influences affecting parental behaviour.

The rest of the thesis is organized as follows. Chapters 2 and 3 relate to the first research question. In particular, chapter 2 focuses on the urban-rural gap in cognitive development observed in Peru and already documented above. Chapter 3 is centred on the empirical strategy proposed to decompose this gap between early childhood and school influences. The chapter analyses how this strategy compares to other linear decomposition methods and discusses in which ways it is superior to the empirical strategies employed so far in the literature. Research questions two and three are examined in Chapter 4, which focuses on the effects of a recently launched parenting programme in rural Peru. Chapter 5 analyses the fourth research question. It explores which caregiver characteristics influence the effect of this programme and proposes a way of thinking about the mechanisms behind caregiver behavioural change. Finally, chapter 6 closes with a review of the main results of the thesis and their potential policy implications. I also discuss the main limitations of the analysis and how these translate into opportunities for future research.

Chapter 2

Explaining the Urban-Rural Gap in Cognitive Achievement in Peru: The role of Early Childhood and School Influences

Abstract

In Peru, students attending rural schools show extremely poor learning outcomes and obtain results significantly below those of urban schools. Because the process of cognitive skill formation is cumulative, differences in initial endowments, early environments and influences occurring later at home and at school can all play a role in shaping these gaps. Previous analyses using Peruvian data report results that favour the role of household characteristics over that of schools or community level variables. The possibility that schools play only a minor role in shaping achievement gaps seems to be supported by the fact that these differences emerge before children enter school and remain almost unchanged throughout school age. This analysis aims at measuring the contribution of school and early childhood influences to the difference in cognitive development observed, at age 8, between urban and rural children in Peru. It contributes new evidence based on an unusually rich dataset and a decomposition strategy less prone to biases than those used so far in the literature. Results indicate that school influences occurring between ages 6 and 8 account for a significant share of urban/rural cognitive skill gap (around 35% and no less than 28%). The share attributable to differences in the early childhood environment is no larger than 50%. Because the gap in cognitive skill was already present at age 5, these results imply that observable differences between rural and urban schools are serving to sustain the inequalities in cognitive development that emerged early in the life of these children.

2.1 Introduction and motivation

In Peru, students attending rural schools show extremely poor learning outcomes and obtain results significantly below those of students in urban schools. National student evaluations reveal a significant and persistent 30 percentage-point gap between the proportion of urban and rural students that attain satisfactory results in reading comprehension (see Figure 2.1 at the end of this chapter). Poverty incidence is also appreciably higher in rural areas, where 46

per cent of the population lives below the national poverty line, versus only 15.3 per cent in urban areas (INEI, 2015).

Learning outcomes measured through reading comprehension and mathematics tests can be understood as a reflection of a broader process of cognitive skill formation. This process is cumulative and, therefore, achievement gaps such as the one documented above must be understood as the result of the effects of all relevant influences on cognitive development until the time at which outcomes are measured. In principle, differences in initial endowments, early environments and influences exerted later both at home and at school can all play a role in shaping these gaps.

In this regard, longitudinal evidence discussed in Chapter 1 reveals that cognitive achievement gaps observed between urban and rural primary school students (around age 8) are already present before these children enter school (age 5) (Figure 2.2 at the end of this chapter illustrates this claim). Moreover, the fact that pre-school gaps have a size similar to those observed later during primary education apparently suggests that school influences play only a subsidiary role, while differences in early childhood environments are much more important. This type of evidence has been recently documented for several developing countries and has been used by some authors to highlight the importance of concentrating investment efforts during early childhood (Heckman, 2006; Schady et al., 2014).

In the education literature, considerable debate surrounds the measurement of the contributions of home background and school quality to children's cognitive development and learning outcomes. Pioneering results were presented in the Coleman Report (Coleman et al., 1966) and the Plowden Report (Peaker, 1971) which used data from the United States and England, respectively. These results suggested that influences occurring prior to school were more significant than school resources when explaining differences in student

outcomes. Two widely cited articles published several years later employed data from developed and developing countries and found that the relationship between school resources and learning outcomes was negatively correlated with national per capita income (Heyneman and Loxley, 1982, 1983). These authors found that school characteristics mattered more than family backgrounds in low income countries. More recent studies have failed to confirm the negative association between national income and the proportion of variance in learning outcomes explained by school characteristics, and have questioned the extent to which, when fully accounting for backgrounds, school resources are effectively translated into improved cognitive outcomes (Baker et al., 2002; Hanushek and Luque, 2003).

Several empirical efforts have tried to address this debate by decomposing learning outcome gaps between children of dissimilar backgrounds into different groups or categories of variables. These studies have distinguished, primarily, between “family” and “school” influences and, in Latin America, have focused on learning gaps between indigenous and non-indigenous children (Arteaga and Glewwe, 2014; Hernandez-Zavala et al., 2006; McEwan, 2004; McEwan and Trowbridge, 2007).

Results from the four studies cited above are not robust and therefore fail to clarify the role played by schools in shaping existing learning outcome gaps and how it compares to the role played by influences exerted earlier in the lives of children. For example, estimates can vary between 17% (Hernandez-Zavala et al., 2006) and 70% (McEwan and Trowbridge, 2007) when it comes to measures of the contribution of schools to the difference in learning outcomes between indigenous and non-indigenous children in Guatemala.

Two of the four cited studies employ data from Peru. In Hernandez-Zavala et al. (2006), the authors considered test scores in mathematics and language obtained by children between

ages 8 and 10 in Guatemala, Mexico and Peru. They measured the contribution of “family and child inputs” (such as parental education, household assets, presence of books at home, and preschool attendance) and “school inputs” (such as teacher experience, access to textbooks, and pupil-teacher ratios) to the indigenous-non-indigenous achievement gap. They found that, in general, family variables contributed more than school variables (23%-33% vs. 17%-23% in Guatemala; 67%-75% vs. 0% in Mexico; 38%-41% vs. 25%-32% in Peru).

In a more recent analysis, Arteaga and Glewwe (2014) used test scores from the Peabody Picture Vocabulary Test (PPVT) and a mathematics test taken by the same group of children at ages 5 and 8. They measured the contribution of “household and child characteristics” (such as household expenditure, parental education, if parents helped with homework and children’s nutritional status) and “community characteristics” (captured through community fixed effects) to the indigenous-non-indigenous achievement gap present at both ages. They found that by the time children were enrolled in school (age 8), household and child characteristics accounted for a significant portion of the achievement gap (around 80%) while community level variables appear to play only a minor role.

Determining which influences and environments play a larger role in shaping achievement gaps can be informative for policy action aimed at overcoming developmental setbacks affecting children from disadvantaged backgrounds. Strong correlations between school characteristics and family backgrounds are a feature of many developing countries and Peru is no exception (Beltran and Seinfeld, 2012; Cueto et al., 2013; Cueto, Guerrero, et al., 2014; Cueto, Leon, et al., 2014). This potential for mutually reinforcing forms of disadvantage means that such an enquiry is especially important from an equity perspective,

in relation to the role of schooling in mitigating the effects of disadvantage caused by home backgrounds.

In this regard, the few results available for Peru and reported above tend to favour the role of household characteristics over that of schools or community level variables. These results seem supported by the fact that gaps in cognitive achievement between children from disadvantaged and more affluent backgrounds emerge before they enter school and persist over time, as documented above. Before concluding that school influences play a subsidiary role when compared to early parental investments and home environments, however, we need to acknowledge some empirical challenges that have not yet been fully addressed by the literature.

In particular, one faces two types of bias when attempting an exercise such as the ones reviewed above. The first is related to the presence of bias in the estimates of individual effects of the determinants of skill. The second can emerge when assigning these determinants to different categories or groups of variables and is caused by the use of rules that end up assigning part of the contribution of one category to another. This second empirical challenge has not yet been addressed in the literature.

Both types of bias have their origin in the presence of unobserved influences. In the first case, strong correlations between school characteristics and family backgrounds complicate the estimation of individual effects as the presence of confounders is more likely. Due to the cumulative nature of the cognitive skill formation process, one needs a particularly rich database (accounting for past and contemporaneous influences) to be able to minimize the risk of omitted variable bias. In this regard, three of the four studies cited above relied only on cross-sectional datasets that lacked information on past influences affecting the

acquisition of skill. This means their estimates of the individual effects of inputs are particularly susceptible to the presence of omitted variable biases.

The risk of incurring in the second type of bias arises from the use of variables that control for omitted influences and the need to assign these controls to a certain category as part of the decomposition exercise. As documented above, is common practice to assign these controls to the category containing “family” influences as these variables typically comprise household or family characteristics (such as family income or parental education). This entails the risk of overstating the contribution of “family” or “household” influences as these variables can also control for omitted inputs that belong to the school environment. As will be discussed later in this chapter, it is useful to distinguish between skill inputs (i.e. variables that have a direct effect on skill such as educational materials provided at home) and skill input determinants (such as family income) when devising a decomposition strategy, in order to minimise the risk of incurring in this second type of bias.

From the discussion above it is clear that there is still an unresolved question regarding the relative importance of schools and early childhood influences for unequal developmental outcomes among children living in the developing world. In the Peruvian context, a prominent way in which these inequalities materialize is in terms of differences in learning outcomes and cognitive achievement between children living in urban and rural areas. As documented in Chapter 1, in Peru, the urban/rural divide constitutes an important axis of inequality in child development.

Considering the above, this study aims to advance the literature on skill formation and child development by measuring the relative importance of school and early childhood influences for the gap in cognitive skill between urban and rural school-age children in Peru. This objective will be addressed using: (i) a comprehensive dataset that includes longitudinal

information on cognitive skills and a wide range of child and household characteristics, combined with detailed information at the school level; and (ii) a decomposition strategy which acknowledges the difference and relations between skill inputs and skill input determinants. The first feature allows one to use time variation in cognitive achievement to identify the individual effect of observed influences. Regarding the second feature, this is the first study to use a decomposition strategy that explicitly acknowledges the role of skill input determinants in the estimation of a production function of skill. This makes the analysis less prone to the second type of bias described above.

The rest of the chapter is organized as follows. Section 2.2 summarises key issues of access, learning outcomes and inequalities in the Peruvian system of basic education. This evidence is complemented with urban/rural gaps in specific school influences when I present the data sources and descriptive statistics in section 2.4.3. Section 2.3 presents the “skill formation technology” and further discusses the two types of bias that can affect an empirical exercise aimed at decomposing an achievement gap. Section 2.4 introduces the decomposition strategy. The results are presented and discussed in section 2.5, and section 2.6 closes with some final remarks.

2.2 Educational inequality in Peru

Peru has achieved high levels of enrolment in education, especially at primary level, where the net enrolment rate stood at 94% in 2011 (WB, 2014). Both primary and secondary education, which cover the ages of 6 to 16, are compulsory and free of tuition fees in public schools, following the 2003 General Education Law. At least one year of pre-schooling (which is compulsory), provided by the state and by private providers, is accessed by a vast majority of children and coverage is being extended to children from the age of 3. The gross enrolment ratio for pre-schooling overall was 78% in 2012 (WB, 2014). Despite high

enrolment from a comparatively young age, however, available evidence suggests that learning levels in Peru are generally low, as well as being highly unequal.

In international testing exercises, Peru's ranking is among the lowest among participating countries in Latin America and is below the levels found in comparable middle-income countries (WB, 2007). In the recent Programme for International Student Assessment (PISA) 2012 exercise, Peru ranked last in mathematics, reading and science out of 65 countries (OECD, 2013).

Economic conditions alone do not explain poor performance, drawing attention to issues of school quality, especially in the public school system. Within this system, the World Bank highlights issues related to weaknesses of management and accountability, pedagogical development with regard to teaching in disadvantaged areas, and issues of linguistic and cultural discrimination (WB, 2007).

In addition to general issues of school quality, part of the explanation for low overall levels of performance lies in the "long tail" of very poorly performing pupils and schools found in a highly unequal system. In comparative terms, Peru's test scores are among the most unequally distributed among all countries for which data are available, with levels of inequality being similar to those in South Africa, which endured decades of enforced discrimination, although there is some evidence of improvement over time (Crouch et al., 2009).

Inequalities in learning outcomes are accompanied by an unequal distribution of educational inputs. While public expenditure on basic education is not especially regressive,¹⁵ the tendency of wealthier parents to choose private schools and highly unequal levels of private

¹⁵ As noted in Crouch et al. (2009), this is not a result of policy but of private choice: wealthier families self-select out of the public system.

spending on education contribute to a regressive distribution of infrastructural and supply inputs across the sector. Recent studies show significant gaps in “opportunities to learn” (OTL)¹⁶ between more and less advantaged pupils in Peru, along axes of language and indigenusness, socio-economic status and urban/rural residence (Cueto, Guerrero, et al., 2014; Cueto, Leon, et al., 2014).

In Table 2.1, I provide additional evidence regarding inequalities in educational inputs at the school level, using the data employed in the decomposition exercises that follow. Information at the school level was collected in 2011 as part of the Young Lives Study in Peru. This survey contains rich information on school and teacher characteristics from a random sample of schools attended, at age 10, by the younger cohort of children followed by this study.

As discussed above, disparities are not especially wide in terms of basic education access. In fact, average years of schooling at the age of 8 are between 2.2 and 2.6 years if we compare the lowest and highest wealth quartiles.¹⁷ There is, however, a strong correlation between children’s wealth and the quality of their school. In particular, access to basic services and educational infrastructure, teacher qualifications, the use of learning material and curriculum coverage all exhibit positive and significant wealth gradients. Increases in school quality as one moves up in the wealth distribution are linked to a shift from rural to urban settlements and to higher enrolment rates in private schools. Wealthier families can afford to pay more for privately provided educational services and also have access to better public schools in urban areas.

¹⁶ OTL typically reflects curriculum content and coverage. In Cueto, Guerrero, et al. (2014) authors employed four variables to account for OTL: curriculum coverage, hours of mathematics at school per year, quality of teachers’ feedback and the level of cognitive demand of exercises.

¹⁷ Based on a wealth index comprising dwelling characteristics, access to basic services and durable goods consumption.

Importantly, children's "school readiness" (measured by their cognitive achievement prior to commencing primary education) also exhibits a positive and highly significant wealth gradient. As discussed in the preceding section, this combination of mutually reinforcing forms of disadvantage renders the effort of comparing the importance of early childhood and school influences both relevant in terms of policy and challenging in terms of empirical strategy.

Table 2.1

Children's cognitive achievement, years of schooling, and school and teacher characteristics by student's wealth^(a) quartile (Q)

	Q1	Q2	Q3	Q4
Household is urban (%)	0.288	0.509 0.221	0.944 0.656***	0.992 0.704***
Standardized raw PPVT score (age 5)	-0.776	-0.478 0.298***	0.322 1.099***	0.87 1.646***
Years of schooling (age 8)	2.208	2.228 0.02	2.458 0.25**	2.553 0.345***
<i>School infrastructure and organization (%)</i>				
School is private	0.00	0.044 0.044	0.146 0.146**	0.268 0.268***
School has basic services ^(b)	0.120	0.281 0.161*	0.833 0.713***	0.894 0.774***
School has a library	0.256	0.456 0.200***	0.66 0.404***	0.659 0.403***
School has a sports ground or play ground	0.552	0.667 0.115	0.896 0.344***	0.935 0.383***
School is multigrade ^(c)	0.480	0.211 -0.269**	0.069 -0.411***	0.057 -0.423***
<i>Teacher characteristics</i>				
Teacher experience (years)	16.165	18.414 2.249	20.548 4.384**	19.118 2.953
Teachers have a university degree (%)	0.349	0.415 0.066	0.525 0.176**	0.703 0.354***
<i>Material and curriculum (%)</i>				
Teachers use books and workbooks	0.189	0.234 0.045	0.34 0.151***	0.359 0.169***
Teachers use complementary materials	0.177	0.26 0.083***	0.313 0.136***	0.31 0.133***
Curriculum coverage (in depth)	0.456	0.504 0.048	0.55 0.095**	0.597 0.141***
Number of obs.	125	114	144	123

Second row for each variable indicates the difference with respect to Q1.

*** p<0.01, ** p<0.05, * p<0.1

(a) Based on a wealth index comprising dwelling characteristics, access to basic services and durable goods consumption.

(b) Basic services comprise water (from a public network or pipe), sanitation (public network connection or a treated cesspool), electricity and telephone connection.

(c) "Multigrade" means that children from different grades receive classes at the same time, in the same room, and from the same teacher.

2.3 Determinants of cognitive skill and empirical challenges for a gap decomposition

2.3.1 The production function of skill

The main objective of this analysis is to estimate the contribution of direct influences occurring at school and of direct influences exerted during early childhood, to the difference in cognitive skill between urban and rural school-age children in Peru.¹⁸

Consistent with this objective, let us divide the relevant phase of child development into two time periods. The first begins when the child is born and finishes at age 5, that is, when the child is ready to start the basic education cycle and enrol in primary school. The second period corresponds to the time when the child remains within primary school age, which is usually between ages 6 and 11.

The production function of cognitive skill provides a good starting point to define its determinants. I assume that skill exhibited by child i at the end of period 2 (A_{i2}) is a function of contemporaneous and past direct influences affecting the child. This is consistent with the notion that skill formation is a cumulative process. Formally:

$$A_{i2} = A_2(HI_{i2}, HI_{i1}, SI_{i2}, SY_{i2}, h_{i2}, h_{i1}, f_i, \mu_{i0}) \quad (2.1)$$

where HI_{i1} are educational inputs provided during early childhood (period 1); HI_{i2} are educational inputs provided at home during period 2; SI_{i2} are educational inputs provided at the school where the child is enrolled during period 2; SY_{i2} are years of schooling attained during period 2; h_{it} indicates the child's health status during period t ; f_i captures predetermined direct influences; and μ_{i0} is the child's innate ability.

¹⁸ This objective is consistent with the main virtue of a decomposition exercise. As discussed in Fortin et al. (2011), decomposition methods will not serve to uncover the mechanisms underlying the relationship between factors and outcomes, but are useful to indicate which factors are quantitatively more important than others

Importantly, equation (2.1) denotes a structural relationship between skill and those variables that have a *direct* effect on it. These variables will reflect the environment surrounding the child (characterizing activities, materials and individuals), as well as child characteristics that influence directly the acquisition of skill. As stressed in Glewwe and Miguel (2008), all the variables in the production function should affect skill directly, and all the variables with a direct effect should be included in this function. For this analysis, I further classify these direct influences as inputs (if they are determined by families' choices during the period under analysis) or as predetermined (if they are outside the current choice set of families). The arguments in this production function are similar to those proposed in Glewwe and Miguel (2008) except for the presence of f_i . This formulation, thus, allows for predetermined child and parental characteristics (e.g. parental education) to have a direct influence on skill.

According to this formulation, differences in terms of cognitive achievement between children living in rural and urban areas should be related to differences between these geographical domains in terms of the direct influences indicated in (2.1).¹⁹ Complete and accurate data on skills and on all these direct influences would allow one to: (i) estimate the structural relationship (i.e. estimate the individual effect of each direct influence); and (ii) estimate the contribution to the urban-rural gap in skill of different groups or categories of influences.

To see this, let us define the urban-rural gap in cognitive skill as the difference in average skill between urban and rural children by the end of period 2 ($E(A_{i2}|U) - E(A_{i2}|R)$) and consider a linear version of (2.1) such as:

¹⁹ Notice that I am assuming that the functional form and parameter values of the production function of skill are the same for urban and rural children. One way of interpreting this is that I am assuming there are no biological differences between children living in these two geographical domains that can affect their learning process.

$$A_{i2} = HI'_{i2}\gamma_1 + HI'_{i1}\gamma_2 + SI'_{i2}\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f'_i\lambda_{(2)} + \mu_{i0}\beta_{(2)} \quad (2.2)$$

Accordingly, the contribution of school influences (CS_2) and of inputs provided during early childhood (CE_2) to the urban-rural gap can be defined as follows:²⁰

$$CS_2 = (E(SI_{i2}|U) - E(SI_{i2}|R))' \phi_1 + (E(SY_{i2}|U) - E(SY_{i2}|R)) \alpha_1 \quad (2.3)$$

$$CE_2 = (E(HI_{i1}|U) - E(HI_{i1}|R))' \gamma_2 + (E(h_{i1}|U) - E(h_{i1}|R))' \varphi_2 \quad (2.4)$$

Sample means of SI_{i2} , SY_{i2} , HI_{i1} and h_{i1} from the urban and rural domains, together with estimates of parameters ϕ_1 , α_1 , γ_2 and φ_2 , could be employed to calculate the contributions expressed in (2.3) and (2.4), and provide an answer to the research question proposed for this analysis.

Unfortunately, complete and accurate information on all the variables considered in (2.1) is seldom available. This is not only because the complete set of relevant direct influences is difficult to define, but also because several of these influences are difficult to measure (e.g. teachers' pedagogical practices) while others remain unobserved (e.g. children's innate ability).

Omission of relevant direct influences can lead to biased estimates of the contribution of different variable categories in two ways. The first is related to the presence of bias in the estimates of production function parameters of skill influences. The second type of bias stems from the use of rules (explicit or implicit) that assign the contribution of variables that belong to one category to another. In terms of the categories proposed for this study, this would mean that at least part of the contribution of the school influences category is being captured by the group referred to the early childhood environment, or vice-versa. In what

²⁰ Here I refer to the contribution of *inputs* provided during early childhood because I am not considering the effect of predetermined direct influences and innate ability.

follows, I further discuss these two types of bias and summarize the approach proposed to mitigate their effects.

2.3.2 Estimating production function parameters

Correlation between included direct influences and those omitted will lead to biased estimates of production function parameters.²¹ This correlation is likely because skill inputs are related through the decision making process of families.

In this regard, there are two aspects related to the identification of production function parameters for a decomposition exercise that are worth considering. First, expressions given in (2.3) and (2.4) may suggest that a decomposition exercise requires a complete account of all relevant inputs and consistent estimates of all input parameters. This analysis, however, is less ambitious as we do not aim at estimating the effect of every variable reflecting an input of skill. What we require is a consistent estimate of the contribution of all the inputs that belong to a certain category, and this can be achieved even under the presence of bias in the estimate of production function parameters of individual inputs. Consider a situation where some of the variables we propose as inputs are picking up the effect of an omitted input that belongs to the same category. For example, some omitted teacher attribute is correlated with an observable school characteristic and they both belong to the “school inputs” category. As shown in Appendix 2.1, this would certainly lead to a biased estimate

²¹ In Fortin et al. (2011), the authors claim that correlation between the error term and covariates can still allow one to obtain consistent estimates of the “unexplained” part of a decomposition, as long as the dependence structure is the same in the two groups under analysis (what they refer to as the ignorability assumption). In this case, however, identification will require more assumptions as I am attempting a decomposition of the “explained” part (i.e. this analysis goes beyond separating the gap between an “explained” and an “unexplained” component). These assumptions will be explicitly listed in the analysis that follows.

of the direct effect of the included school characteristic, but can still allow one to recover a consistent estimate of the contribution of the “school inputs” category.²²

The second aspect worth considering is that equations (2.3) and (2.4) involve several inputs and are expressed in terms of production function parameters. This implies that experimental designs and quasi-experimental techniques are not only impractical but also unsuited to estimating the contributions of interest. Lack of practicality stems from the difficulty of randomly assigning several different treatments (one for each input under analysis) or, if we aim at using an instrumental variable technique, it stems from the difficulty of finding enough sources of exogenous variation.

Equally important is the fact that these techniques are particularly well suited to estimating the *total* effect of a particular input and not to estimating its production function parameter or *direct* effect. Parameters involved in (2.2) indicate a direct effect because they capture the result of shifting a particular influence holding the rest constant. Random assignment to a treatment group or a valid instrumental variable, however, will provide an estimate of this direct effect plus an indirect one which occurs through changes in other inputs (recall that inputs are choice variables and can therefore change in response to a shock to some skill determinant).²³ Experimental and quasi-experimental techniques, thus, would not allow one to recover the contribution of different categories of inputs in terms of the production function parameters as indicated in (2.3) and (2.4).

²² This is possible because the contribution of the omitted input can be expressed as a function of the bias in the estimated individual effect of the observed input and its difference in means. Thus, the contribution of the omitted input will still be assigned to its corresponding category if it is the same as the category of the observed input (see Appendix 2.1).

²³ This means that experimental and quasi-experimental techniques usually recover the parameters of a conditional demand function and not the parameters of a production function. See Glewwe and Miguel (2008) for a discussion of the differences between these two types of functions. In Todd and Wolpin (2003) the authors distinguish between “policy effects” and “technology parameters” and provide some examples.

The discussion above implies that the identification strategy for the production function parameters of observed direct influences will have to be based on employing a rich set of controls and imposing some assumptions on the production function of skill and the behaviour of its determinants. In Todd and Wolpin (2003) and Todd and Wolpin (2007), the authors have already discussed the assumptions required to obtain consistent estimates of production function parameters under different empirical specifications. Here I briefly review the assumptions related to two particular specifications (the hybrid model and the value added model) as they will play an important role in the empirical strategy.

(a) *The hybrid model*

Let us return to the linear production function specified in (2.2) and assume that skill is measured with error through the scores obtained in some cognitive test: $T_{i2} = A_{i2} + \varepsilon_{i2}$. Thus, we obtain:

$$T_{i2} = HI'_{i2}\gamma_1 + HI'_{i1}\gamma_2 + SI'_{i2}\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f'_i\lambda_{(2)} + \mu_{i0}\beta_{(2)} + \varepsilon_{i2} \quad (2.5)$$

As already noted, consistent estimation of the parameters involved in the expression above is problematic because we seldom observe all the relevant direct influences. A common practice is to try to circumvent this problem by substituting missing inputs with their corresponding demand functions, which will typically contain variables reflecting family resources, prices, predetermined direct influences of skill, exogenous environmental characteristics, child's skill and health endowments, and parental preferences.²⁴ This gives rise to what the literature has described as a “hybrid” function (Rosenzweig and Schultz, 1983; Todd and Wolpin, 2007).

²⁴ See Glewwe and Miguel (2008) for a good exposition regarding input demand functions.

To illustrate this, assume there are only two period 1 educational inputs: one that can be observed (HI_{i1}^O) and one that cannot (HI_{i1}^U). Additionally, assume that the demand for the unobserved input can be expressed as a linear function of the predetermined direct influences of skill (f_i) and a vector of exogenous input determinants (z_i), up to an error term (v_i): $HI_{i1}^U = z_i'\delta + f_i'\kappa + v_i$. This enables us to replace the unobserved input for its demand function to obtain:

$$\begin{aligned}
T_{i2} &= HI_{i2}'\gamma_1 + HI_{i1}^O\gamma_2^O + SI_{i2}'\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f_i'\lambda_{(2)} + \\
&\quad + HI_{i1}^U\gamma_2^U + \mu_{i0}\beta_{(2)} + \varepsilon_{i2} \\
&= HI_{i2}'\gamma_1 + HI_{i1}^O\gamma_2^O + SI_{i2}'\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f_i'\lambda_{(2)} + \\
&\quad + (z_i'\delta + f_i'\kappa + v_i)\gamma_2^U + \mu_{i0}\beta_{(2)} + \varepsilon_{i2} \\
&= HI_{i2}'\gamma_1 + HI_{i1}^O\gamma_2^O + SI_{i2}'\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f_i'\pi + \\
&\quad + z_i'\psi + e_{H,i}^{CU}
\end{aligned} \tag{2.6}$$

where $\pi = (\lambda_{(2)} + \kappa\gamma_2^U)$, $\psi = \gamma_2^U\delta$ and $e_{H,i}^{CU} = \gamma_2^U v_i + \mu_{i0}\beta_{(2)} + \varepsilon_{i2}$. As discussed in Todd and Wolpin (2007), there should be no correlation between included inputs and the error term v_i for this strategy to overcome the potential omitted variable bias.

(b) *The value added model*

The expression given in (2.5) corresponds to a “cumulative” model. It is a specification that accounts for all the relevant influences received until the period in which skill is measured. A specification that only accounts for contemporaneous influences and controls for past achievement is typically known as a “value added” model. Formally:

$$T_{i2} = \rho T_{i1} + HI_{i2}'\gamma_1 + SI_{i2}'\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + f_i'\lambda + e_{P,i}^{VA} \tag{2.7}$$

where $e_{p,i}^{VA} = \mu_{i0}\beta + \varepsilon_{i2} - \rho\varepsilon_{i1}$. This is a restricted version of the cumulative specification, with the restriction being that the effects of all period 1 influences decline at the same rate ρ . To see this, notice that period 1 skill can be expressed as:

$$T_{i1} = HI'_{i1}\gamma_1 + h_{i1}\varphi_1 + f'_i\lambda + \mu_{i0}\beta + \varepsilon_{i1} \quad (2.8)$$

so that:

$$\begin{aligned} T_{i2} - \rho T_{i1} = & HI'_{i2}\gamma_1 + HI'_{i1}(\gamma_2 - \rho\gamma_1) + SI'_{i2}\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}(\varphi_2 - \rho\varphi_1) + \\ & f'_i(\lambda_{(2)} - \rho\lambda) + \mu_{i0}(\beta_{(2)} - \rho\beta) + \varepsilon_{i2} - \rho\varepsilon_{i1} \end{aligned} \quad (2.9)$$

For (2.9) to reduce to (2.7) the following must hold: $\gamma_2 = \rho\gamma_1$, $\varphi_2 = \rho\varphi_1$, $\lambda_{(2)} = (1 + \rho)\lambda$, and $\beta_{(2)} = (1 + \rho)\beta$. This means that the rate of decay of the effects of influences occurring in period 1 is the same for all influences and equal to ρ .²⁵

Consistent estimation of the effect of contemporaneous influences requires absence of correlation between the error term $e_{p,i}^{VA}$ and the rest of the right hand variables in (2.7). This assumption is particularly problematic for lagged measured skill (T_{i1}) because it will be correlated with its own measurement error unless, as pointed out by Todd and Wolpin (2007), we assume that measurement error is serially correlated with an autoregressive parameter equal to the rate of decay of the effect of inputs. The presence of innate ability in the error term also constitutes a potential source of bias for the OLS estimate of the persistence parameter (ρ) and of input parameters.

²⁵ Notice that I am allowing predetermined direct influences and innate ability to exert an effect every period equal to λ and β , respectively. If these influences decay at a rate ρ , it means that, by period 2, the cumulative effect will be given by $\lambda_{(2)} = \lambda + \rho\lambda$ and $\beta_{(2)} = \beta + \rho\beta$, which correspond to the conditions given above.

An interesting feature of the potential biases affecting the estimation of the persistence parameter is that they operate in opposite directions. As discussed by Andrabi et al. (2011), innate ability will produce a positive bias while measurement error will cause the typical attenuation effect. If we combine this result with the fact that lagged achievement is partially controlling for innate ability, the overall bias affecting ordinary least square (OLS) estimates of contemporaneous input effects can be small.²⁶

The inclusion of past inputs in a value added model allows one to partially relax the assumption that the effects of past inputs share the same rate of decay. This yields:

$$T_{i2} = \rho T_{i1} + HI'_{i2} \gamma_1 + HI'_{i1} \tilde{\gamma}_2^O + SI'_{i2} \phi_1 + SY_{i2} \alpha_1 + h_{i2} \varphi_1 + h_{i1} \tilde{\varphi}_2 + f'_i \lambda + e_{p,i}^{VAP} \quad (2.10)$$

In an specification like (2.10) (which will be referred to as the “value added plus” specification following Todd and Wolpin (2007)), the effects of included past inputs are allowed to have different rates of decay, which are accommodated through the parameters contained in $\tilde{\gamma}_2^O$ and $\tilde{\varphi}_2$. It should be noted that these parameters are not the same as γ_2^O and φ_2 given in (2.6). In fact, inspection of (2.9), above, reveals that $\tilde{\gamma}_2^O = \gamma_2^O - \rho \gamma_1^O$ and $\tilde{\varphi}_2 = \varphi_2 - \rho \varphi_1$. This provides a simple test for the restriction that the effects of all period 1 influences decline at the same rate ρ , which consists in evaluating the significance of past inputs in a value added plus specification. It also shows that in case the restriction fails to hold, parameters $\tilde{\gamma}_2^O$ and $\tilde{\varphi}_2$ will accommodate the differences.

2.3.3 Assigning the contribution of input determinants

Recall I argued that the omission of relevant direct influences can lead to biased estimates of the contribution of different variable categories in two ways. The first was discussed in the

²⁶ This result is proposed by Andrabi et al. (2011) to explain why OLS estimates of the private-school skill premium in Pakistan provided by a valued added model turned out similar to those obtained after applying item response theory and dynamic panel data methods to mitigate the effects of measurement error and observed innate ability, respectively.

preceding section and is related to the presence of bias in the estimates of production function parameters. The second type of bias has been overlooked so far in the literature. It also has its roots in the presence of omitted direct influences and has to do with the use of predetermined family and child characteristics to control for unobservable inputs. Predetermined family and child characteristics can be proposed for this as long as they refer to skill input determinants. In other words, as long as they are relevant arguments in the demand function of skill inputs. As explained above, doing this configures a “hybrid” model.

Input determinants can control for several omitted inputs, so one needs to be especially careful when assigning these controls to a particular variable category. A good example of this is family income. Family income is not an input into skill formation, but an input determinant. This means it has only an indirect effect on skill through the inputs that are sensitive to families’ budget constraints. If family income has a role determining the quantity or quality of inputs received during early childhood and also later at school, and there are unobservable influences related to both environments, it would not be appropriate to assign the contribution of this control exclusively to either of them. Doing so would lead to a biased estimate of the contribution of the category hosting family income.

Based on the above, in this analysis I will assign the contribution of predetermined direct influences and other input determinants to a special category hosting omitted inputs. This is consistent with the notion that family resources and preferences, as well as local educational supply characteristics, play a role in determining not only the educational environment during early childhood (HI_{i1}), but also the number of years of schooling (SY_{i2}) and the educational inputs provided at school (SI_{i2}) during period 2. As described in Section 2.2, a key feature of the Peruvian education system is that school characteristics are highly

heterogeneous and strongly correlated with children's socioeconomic status. The use of a special category to host omitted inputs in general will also prevent one from making strong assumptions regarding the nature of unobservable influences.

2.4 Decomposition strategy

2.4.1 Empirical specifications and assumptions

The four empirical specifications that stem from combining the hybrid and the value added models described above are summarized in Table 2.2. Specifications labelled under the category “production function” only consider influences (both inputs and predetermined) with a direct effect. Those labelled as “hybrid” also consider influences with an indirect effect (i.e. exogenous input determinants). Cumulative specifications only consider the entire history of influences while those grouped under the term “value added plus” also control for past measured skill.

Table 2.2
Empirical specifications

	<i>Cumulative</i>	<i>Value added plus</i>
Production function	$T_{i2} = HI_{i2}^{O'}\gamma_1 + HI_{i1}^{O'}\gamma_2 + SI_{i2}^{O'}\phi_1$ $+ SY_{i2}\alpha_1 + h_{i2}\varphi_1$ $+ h_{i1}\varphi_2 + f_i'\lambda_{(2)}$ $+ e_{P,i}^{CU}$	$T_{i2} = \rho T_{i1} + HI_{i2}^{O'}\gamma_1 + HI_{i1}^{O'}\tilde{\gamma}_2$ $+ SI_{i2}^{O'}\phi_1 + SY_{i2}\alpha_1$ $+ h_{i2}\varphi_1 + h_{i1}\tilde{\varphi}_2 + f_i'\lambda$ $+ e_{P,i}^{VAP}$
Hybrid	$T_{i2} = HI_{i2}^{O'}\gamma_1 + HI_{i1}^{O'}\gamma_2$ $+ SI_{i2}^{O'}\phi_1 + SY_{i2}\alpha_1$ $+ h_{i2}\varphi_1 + h_{i1}\varphi_2$ $+ f_i'\pi + z_i'\psi + e_{H,i}^{CU}$	$T_{i2} = \rho T_{i1} + HI_{i2}^{O'}\gamma_1 + HI_{i1}^{O'}\tilde{\gamma}_2$ $+ SI_{i2}^{O'}\phi_1 + SY_{i2}\alpha_1$ $+ h_{i2}\varphi_1 + h_{i1}\tilde{\varphi}_2 + f_i'\tilde{\pi}$ $+ z_i'\tilde{\psi} + e_{H,i}^{VAP}$

The elements contained in the error terms of these specifications play a crucial role in determining the existence of bias in an OLS estimate of the production function parameters.

Let us consider a minimum set of assumptions that will serve to impose some structure on these error terms and to identify the production function parameters of interest.

- (i) The production function of skill can be expressed as a linear function of inputs and predetermined direct influences as follows:

$$A_{i2} = HI'_{i2}\gamma_1 + HI'_{i1}\gamma_2 + SI'_{i2}\phi_1 + SY_{i2}\alpha_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f'_i\lambda_{(2)} + \mu_{i0}\beta_{(2)}$$

$$A_{i1} = HI'_{i1}\gamma_1 + h_{i1}\varphi_1 + f'_i\lambda + \mu_{i0}\beta$$

This linearity assumption will allow one to directly compare the decomposition strategy proposed here to the Oaxaca-Blinder technique commonly employed in the literature. This will be further discussed in Chapter 3. This assumption comes at the cost of ignoring potential complementarities between inputs. One type of complementarity that could be of particular relevance is what the literature refers to as “dynamic complementarity” (Cunha and Heckman, 2007). That is, the relation between the effects of contemporaneous inputs and past cognitive achievement. One way of partially relaxing the linearity assumption and allowing for this type of complementarity is by adding interactions between contemporaneous inputs and past cognitive skills. This will be tested in the next chapter.

- (ii) Skill is measured with random error: $T_{it} = A_{it} + \varepsilon_{it}; E(\varepsilon_{it}) = 0$, $Cov(A_{it}, \varepsilon_{it}) = 0; t = 1, 2$.
- (iii) Innate ability (μ_{i0}) cannot be observed. Some period 1 and period 2 inputs also remain unobserved and are contained in vectors UI_{i1} and UI_{i2} , respectively. Years of schooling, health inputs and predetermined direct influences are observed.²⁷

²⁷ Input variables such as years of schooling and exogenous input determinants such as parental education could also be measured with error. This also applies to variables reflecting school quality (school inputs). This is not explicitly addressed in the formulations that follow. It is reasonable to argue, however, that this analysis should not be especially susceptible to the potential attenuation biases as it does not aim at estimating the effect of a particular input of skill but the *relative* importance (or contribution to a cognitive skill gap) of variable categories that involve several inputs.

- (iv) Skill inputs are determined by parents' choices and it is possible to characterize their demand equations as a linear function of predetermined direct influences (f_i) and other exogenous input determinants (z_i) up to a random error (v_i).
- (v) The effects of unobserved period 1 influences and predetermined direct influences decay at a rate equal to ρ .

Assumptions (i), (ii) and (iii) imply that the error term of the production function-cumulative model is given by:

$$e_{P,i}^{CU} = UI'_{i1}\tau_2 + UI'_{i2}\tau_1 + \mu_{i0}\beta_{(2)} + \varepsilon_{i2} \quad (2.11)$$

where τ_2 and τ_1 contain the production function parameters of period 1 and period 2 unobserved inputs.

In addition, assumption (iv) allows one to express the demand functions of unobserved inputs as follows: $UI_{i1} = \delta_1 z_i + \kappa_1 f_i + v_{i1}$ and $UI_{i2} = \delta_2 z_i + \kappa_2 f_i + v_{i2}$.²⁸ This implies that the error term of the hybrid-cumulative model is given by:

$$e_{H,i}^{CU} = v'_{i1}\tau_2 + v'_{i2}\tau_1 + \mu_{i0}\beta_{(2)} + \varepsilon_{i2} \quad (2.12)$$

and that $\pi = \lambda_{(2)} + \kappa'_1\tau_2 + \kappa'_2\tau_1$ and $\psi = \delta'_1\tau_2 + \delta'_2\tau_1$.

Finally, assumption (v) implies that the error terms of the two value added plus specifications can be expressed as:

$$e_{P,i}^{VAP} = UI'_{i2}\tau_1 + \mu_{i0}(\beta_{(2)} - \rho\beta) + \varepsilon_{i2} - \rho\varepsilon_{i1} \quad (2.13)$$

$$e_{H,i}^{VAP} = v'_{i2}\tau_1 + \mu_{i0}(\beta_{(2)} - \rho\beta) + \varepsilon_{i2} - \rho\varepsilon_{i1} \quad (2.14)$$

and that $\tilde{\pi} = \lambda + \kappa'_2\tau_1$ and $\tilde{\psi} = \delta'_2\tau_1$.

²⁸ Notice that UI_{i1} and UI_{i2} are vectors. Accordingly, each row in the matrices δ_1 and κ_1 , and δ_2 and κ_2 , contain the parameters of the demand functions of period 1 and period 2 unobserved inputs, respectively.

Following the discussion presented in section 2.3, the estimation of the production function parameters will be based on the specification that provides the richest set of controls. This is the case of the hybrid-value added plus model. The reasons for preferring this specification can be summarized as follows: (i) it controls for omitted period 1 inputs; (ii) it controls for the effect of innate ability accumulated during period 1; and (iii) the introduction of exogenous input determinants allows one to control for omitted period 2 inputs and mitigate omitted variable bias as long as the error terms of their demand functions are not correlated with observed inputs.

Among these features, being able to control for past (period 1) omitted direct influences is particularly important for obtaining a consistent estimate of the contribution of school inputs. This is because we know that this input category is relevant only in period 2, so correlation between any omitted period 1 influence and included school inputs would lead to a biased estimate of the contribution of school influences. As discussed earlier, omission of a period 2 input will not necessarily bias the estimate of the contribution of school influences if the omitted input belongs to the school environment.

It is also worth noticing that a value added model would be able to fully control for innate ability under the assumption that it only affects initial conditions (i.e. that it has a once-and-for-all effect on skill). As in Andrabi et al. (2011), I have adopted here the more realistic assumption that innate ability affects the learning process in every period. As a consequence, the error terms of both value added models include the term $\mu_{i0}(\beta_{(2)} - \rho\beta)$. This captures the contemporaneous effect of innate ability on skill, and its presence can be of concern for the purpose of identifying the production function parameters of included influences and lagged skill.

In this regard, empirical results presented in several recent studies corroborate that value added models can provide reliable estimates of the individual effects of skill inputs. These studies are reviewed in Singh (2015). They show that value added models such as the one presented in (2.7) outperform other empirical strategies when recovering teacher effects from simulated data (Guarino et al., 2012), and provide the same results as experimental and quasi-experimental methods used to identify school or teacher effects (Deming et al. (2014) Kane et al. (2013), among others). Moreover, value added estimates given in Singh (2015) for the effect of private school enrolment on the achievement of rural children were also found to be similar to the results provided by an experimental exercise carried out in the same region of India (Muralidharan and Sundararaman, 2013).²⁹

These results indicate that lagged achievement is a sufficient statistic to control for assignment mechanisms that correlate with innate ability. In terms of the error structures given above, this implies that: (i) correlation between the contemporaneous effect of innate ability $\left(\mu_{i0}(\beta_{(2)} - \rho\beta)\right)$ and included inputs is not significant;³⁰ and (ii) that potential biases affecting the estimate of ρ (through correlation between lagged test scores and innate ability or measurement error) do not transpire into the estimates of the individual effects of inputs (a point already made by Andrabi et al. (2011) and highlighted again by Singh (2015)).

²⁹ It should be noted that all the empirical exercises cited in this paragraph aimed at obtaining the total (or “policy”) effect of being assigned to a certain type of school input. Nevertheless, assumptions in terms of absence of correlation between the contemporaneous effect of innate ability and included inputs and lagged skill are the same as those required in a value added model aimed at estimating the direct effects (or production function parameters) of inputs. The difference relies on the controls used in the value added specification. Total effects require that one control only for the exogenous determinants of excluded inputs, which correspond to the arguments of their conditional demand functions (see Glewwe and Miguel (2008)). The estimation of production function parameters requires one to control for the rest of relevant inputs. This is consistent with the variables included in the empirical specifications given in Table 2.2, and the use of the hybrid specification as a way of controlling for potentially omitted period 2 inputs.

³⁰ It is reasonable to assume that parents make decisions regarding inputs based on the signals of skill they can observe (e.g. performance at preschool). This information can be summarized by lagged test scores which are already being controlled for in a value added model.

Correlation between lagged test scores and the contemporaneous effect of innate ability and between lagged test scores and measurement error are a potential source of bias for the estimation of the persistence parameter (ρ). To evaluate this, I will also consider the results provided by a value added model after removing the correlation between lagged test scores and measurement error.³¹ Because the two potential biases operate in different directions, one can think of this estimation as providing an upper bound for the contribution of past influences (only the possibility of a positive bias remains in our estimate of ρ) and a lower bound for the contribution of the rest of inputs.³²

The nature of the available data also plays a role in the process of narrowing down a preferred empirical specification. For example, a dataset with a richer set of period 2 influences than period 1 influences (like the one described in section 2.4.3 below) makes a stronger case for value added plus specification.

Preference for the hybrid-value added plus specification is based on the performance of value added models reported in previous studies and on comparing the error structures given in equations (2.11) - (2.14). To provide more evidence in support of this specification choice, the empirical strategy will comprise a comparison of the results obtained in terms of decompositions and parameter estimates across all four specifications. The objective is to confirm if these estimates behave in a manner consistent with the error structures given above.³³

³¹ This can be achieved by instrumenting lagged test scores (T_{i1}) with another measure of cognitive skill. The objective is to induce variation in T_{i1} uncorrelated with its own random measurement error (ε_{i1}). Notice this rests on the assumption of no correlation between the measurement errors affecting the two test scores.

³² Notice that contemporaneous inputs could also be measured with error. The presence of an attenuation bias in the estimated coefficient of these inputs, however, would still allow one to present these results as a lower bound or conservative estimate of the contribution of school inputs to the cognitive skill gap.

³³ It is worth noticing that a dynamic panel data approach to difference out innate ability from the error term of the value added specification is not applicable in this case. This is because I lack baseline (beginning of period 1) test scores and also because it would not allow one to recover parameter estimates for time invariant influences. In Andrabi et al. (2011), the authors proposed a dynamic panel data strategy to identify the private

2.4.2 Urban/rural gap decompositions

The variables involved in each of the four specifications given in Table 2.2 will be arranged in different groups or categories. The empirical goal is to measure the contribution of each category to the estimated period 2 urban-rural gap in cognitive skill, which is given by $\bar{T}_{U2} - \bar{T}_{R2}$ (where an upper bar denotes the sample mean of the variable). This contribution is given by the weighted average of the urban-rural differences in the variables assigned to the category, with weights given by the estimated values of their production function parameters.

If the specification includes an urban/rural group indicator, the sum of the contributions of all categories will equal the estimated urban/rural gap in cognitive skill. This is because the group indicator ensures that the regression line passes through the means of both groups. Importantly, an urban/rural indicator cannot be regarded as a direct influence but can have a role among input determinants.³⁴ Accordingly, the inclusion of this indicator will be possible in the hybrid models but not in the production function specifications.

The categories proposed for this analysis have to allow one to identify the contribution of influences originating at school and of influences that occurred during early childhood (period 1). Categories must also be consistent with the role played by specific variables in the production function of skill. As already discussed, one needs to be especially careful when assigning the contribution of variables that control for omitted inputs. This becomes clearer if we examine the structure behind the parameters associated to exogenous input determinants (z_i) and predetermined direct influences (f_i) in the hybrid models.

school premium in Pakistan and relied on students switching schools to recover this parameter. It is worth recalling that they found estimates very similar to those provided by a value added model.

³⁴ This point is further developed in the next chapter.

Notice that in the hybrid-cumulative specification, exogenous input determinants are related to a parameter vector (ψ), which is a function of two sets of parameters: (i) those of exogenous input determinants in the demand functions of omitted inputs (δ_1 and δ_2); and (ii) those of omitted inputs in the production function of skill (τ_2 and τ_1). In fact, $\psi = \delta_1' \tau_2 + \delta_2' \tau_1$.

Predetermined direct influences have both a direct and an indirect effect, the latter of which operates through the demand for omitted inputs. Their parameters (contained in vector π) are a function of three sets of parameters: (i) the cumulative effect of predetermined direct influences in the production function of skill ($\lambda_{(2)}$); (ii) the parameters of predetermined direct influences in the demand functions of omitted inputs (κ_1 and κ_2); and (iii) the parameters of omitted inputs in the production function of skill (τ_2 and τ_1). In fact, $\pi = \lambda_{(2)} + \kappa_1' \tau_2 + \kappa_2' \tau_1$.

Assumption (v), above, implies that a value added specification will effectively remove period 1 unobserved inputs as well as the effect, in period 1, of predetermined direct influences. Accordingly, in the hybrid-value added plus model, elements in $\tilde{\psi}$ and $\tilde{\pi}$ will no longer depend on the production or demand function parameters of period 1 omitted inputs. As indicated above, $\tilde{\pi} = \lambda + \kappa_2' \tau_1$ and $\tilde{\psi} = \delta_2' \tau_1$.

An important implication of the parameter structure described above is that, in absence of further restrictions, it will not be possible to separately identify the direct and indirect effects of predetermined direct influences in a hybrid specification. Strong assumptions are also required to claim that omitted inputs belong only to either the early childhood or school environment.³⁵ As a consequence, for hybrid specifications I will jointly measure the

³⁵ The use of school fixed effects would allow one to claim that period 2 omitted influences belong only to the home environment. School fixed effects, however, can also pick-up the effects of period 1 influences if

contribution of predetermined direct influences and all omitted inputs by allocating the contribution of variables contained in f_i and z_i into a special category. In the particular case of the hybrid-value added plus specification, this joint contribution will consider only period 2 predetermined direct influences and period 2 omitted inputs.

The parameter structure described above also allows for a simple test for omitted inputs by analysing the significance of the contribution of exogenous input determinants in the hybrid specifications. The intuition is that exogenous input determinants will contribute information to the estimation of the test score gap if skill inputs have not been fully accounted for. In fact, rejection of the null hypothesis $(\bar{z}_U - \bar{z}_R)' \psi = 0$ or $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$ in the corresponding hybrid model implies the presence of at least one omitted input.³⁶ More specifically, rejection of the null $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$ in the hybrid-value added plus specification implies the presence of at least one period 2 omitted input.

It is also worth noticing that in the value added specifications considered I am not restricting the effect of observed inputs to decay at the same rate ρ (hence the name “value added plus”). As discussed when presenting the value added plus model in the previous section, despite including period 1 observed inputs, these specifications will not allow one to recover their direct effects. In fact, the contributions of lagged achievement and period 1 inputs will have to be merged into a single category of “past influences”.

children are sorted into schools according to ability. As a consequence, school fixed effects typically provide large estimates of the contribution of school variables to achievement gaps (see, for example, McEwan (2004) or McEwan and Trowbridge (2007)). This potential source of bias is further discussed in Chapter 3.

³⁶ Consider the case of the hybrid-cumulative model. Rejection of the null $(\bar{z}_U - \bar{z}_R)' \psi = 0$ implies that $\psi \neq 0$. Given that $\psi = \delta'_1 \tau_2 + \delta'_2 \tau_1$, $\psi \neq 0$ implies that either $\delta'_1 \tau_2 \neq 0$ or $\delta'_2 \tau_1 \neq 0$. It suffices for one of these two inequalities to hold to conclude that there is at least one omitted input and that z_i is a relevant argument in its demand function (i.e. either $\tau_2 \neq 0$ and $\delta_1 \neq 0$ or $\tau_1 \neq 0$ and $\delta_2 \neq 0$). Notice that failure to reject the hypothesis $(\bar{z}_U - \bar{z}_R)' \psi = 0$ does not directly imply the absence of omitted inputs. For example, parameters in ψ can be zero even if $\tau_2 \neq 0$ and $\tau_1 \neq 0$ (i.e. there are omitted inputs from both periods) if the proposed exogenous input determinants have no role in their demand equations ($\delta_1 = 0$ and $\delta_2 = 0$). This could be the case if the variables considered as predetermined direct influences (f_i) fully characterize the demand for omitted inputs.

The categories related to the four specifications given in Table 2.2 are summarized in Table 2.3. Notice that in case the regression includes an interaction term between two different inputs, one could consider introducing a new category to account for the contribution of concurrent differences in the two inputs to facilitate the interpretation of decomposition results.³⁷

³⁷ Consider, for example, a value added specification such as: $T_{i2} = \rho T_{i1} + \gamma I_{i2} + \beta T_{i1} I_{i2} + \varepsilon_{i2}$. The urban-rural gap in period 2 test scores could be decomposed as follows: $\bar{T}_{U2} - \bar{T}_{R2} = \hat{\rho}(\bar{T}_{U1} - \bar{T}_{R1}) + \hat{\beta}(\bar{T}_{U1}\bar{I}_{R2} - \bar{T}_{R1}\bar{I}_{R2}) + \hat{\gamma}(\bar{I}_{U2} - \bar{I}_{R2}) + \hat{\beta}(\bar{T}_{R1}\bar{I}_{U2} - \bar{T}_{R1}\bar{I}_{R2}) + \hat{\beta}[(\bar{T}_{U1}\bar{I}_{U2} - \bar{T}_{R1}\bar{I}_{R2}) - (\bar{T}_{U1}\bar{I}_{R2} - \bar{T}_{R1}\bar{I}_{R2}) - (\bar{T}_{R1}\bar{I}_{U2} - \bar{T}_{R1}\bar{I}_{R2})]$. The first two terms at the right hand side of the preceding equation would capture the contribution of past skills *alone*, the next two terms would capture the contribution of the contemporaneous input *alone*, and the last term (in brackets) would capture the additional gain obtained from changes in both inputs at the same time.

Table 2.3

Variable categories related to each empirical specification and their contribution to the urban-rural gap in cognitive skill

	<i>Cumulative</i>		<i>Value added plus</i>	
Prod. function	Early childhood and home inputs	$(\overline{HI}_{U2}^0 - \overline{HI}_{R2}^0)' \hat{\gamma}_1 + (\overline{HI}_{U1}^0 - \overline{HI}_{R1}^0)' \hat{\gamma}_2$	Period 2 home inputs	$(\overline{HI}_{U2}^0 - \overline{HI}_{R2}^0)' \hat{\gamma}_1$
	Health inputs	$(\bar{h}_{U2} - \bar{h}_{R2}) \hat{\phi}_1 + (\bar{h}_{U1} - \bar{h}_{R1}) \hat{\phi}_2$	Period 2 health inputs	$(\bar{h}_{U2} - \bar{h}_{R2}) \hat{\phi}_1$
	School inputs	$(\overline{SI}_{U2}^0 - \overline{SI}_{R2}^0)' \hat{\phi}_1 + (\overline{SY}_{U2} - \overline{SY}_{R2})' \hat{\alpha}_1$	School inputs	$(\overline{SI}_{U2}^0 - \overline{SI}_{R2}^0)' \hat{\phi}_1 + (\overline{SY}_{U2} - \overline{SY}_{R2})' \hat{\alpha}_1$
	Predetermined direct influences	$(\bar{f}_U - \bar{f}_R)' \hat{\lambda}_{(2)}$	Period 2 predetermined direct influences	$(\bar{f}_U - \bar{f}_R)' \hat{\lambda}$
	--	--	Past influences	$\hat{\rho}(\bar{T}_{U1} - \bar{T}_{R1}) + (\overline{HI}_{U1}^0 - \overline{HI}_{R1}^0)' \hat{\gamma}_2 + (\bar{h}_{U1} - \bar{h}_{R1}) \hat{\phi}_2$
Hybrid	Early childhood and home inputs	$(\overline{HI}_{U2}^0 - \overline{HI}_{R2}^0)' \hat{\gamma}_1 + (\overline{HI}_{U1}^0 - \overline{HI}_{R1}^0)' \hat{\gamma}_2$	Period 2 home inputs	$(\overline{HI}_{U2}^0 - \overline{HI}_{R2}^0)' \hat{\gamma}_1$
	Health inputs	$(\bar{h}_{U2} - \bar{h}_{R2}) \hat{\phi}_1 + (\bar{h}_{U1} - \bar{h}_{R1}) \hat{\phi}_2$	Period 2 health inputs	$(\bar{h}_{U2} - \bar{h}_{R2}) \hat{\phi}_1$
	School inputs	$(\overline{SI}_{U2}^0 - \overline{SI}_{R2}^0)' \hat{\phi}_1 + (\overline{SY}_{U2} - \overline{SY}_{R2})' \hat{\alpha}_1$	School inputs	$(\overline{SI}_{U2}^0 - \overline{SI}_{R2}^0)' \hat{\phi}_1 + (\overline{SY}_{U2} - \overline{SY}_{R2})' \hat{\alpha}_1$
	Predetermined direct influences and omitted inputs	$(\bar{f}_U - \bar{f}_R)' \hat{\pi} + (\bar{z}_U - \bar{z}_R)' \hat{\psi}$	Period 2 predetermined direct influences and period 2 omitted inputs	$(\bar{f}_U - \bar{f}_R)' \hat{\pi} + (\bar{z}_U - \bar{z}_R)' \hat{\psi}$
	--	--	Past influences	$\hat{\rho}(\bar{T}_{U1} - \bar{T}_{R1}) + (\overline{HI}_{U1}^0 - \overline{HI}_{R1}^0)' \hat{\gamma}_2 + (\bar{h}_{U1} - \bar{h}_{R1}) \hat{\phi}_2$

Finally, it is worth noting that in the hybrid-value added plus specification, the contribution of direct school influences will be given by the category labelled “school inputs” in Table 2.3. The contribution of direct influences that occurred during early childhood will be estimated through the category labelled “past influences”. This category includes the contribution of educational and health inputs provided in period 1, and also the contribution of predetermined direct influences and innate ability that is due to their period 1 effects (which are subsumed in lagged achievement).

2.4.3 Data sources and variables

The data for this analysis will be provided by the Peruvian dataset of the Young Lives Study. Young Lives is an international study of childhood poverty, following 12,000 children in 4 countries (Ethiopia, India, Peru and Vietnam) over 15 years. In particular, it will consider the first three rounds of the child and household surveys, as well as the school survey, focusing on the Younger Cohort. Table 2.4 summarises the basic time structure of these data.

Table 2.4
Time-structure and sample sizes of the relevant Young Lives databases

	<i>Child and household survey</i>			<i>School Survey</i>
	<i>Round 1</i> 2002	<i>Round 2</i> 2006	<i>Round 3</i> 2009	2011
Younger cohort's age (years)	1 (0.5-1.5)	5 (4.5-5.5)	8 (7.5-8.5)	10 (9.5-10.5)
Sample size (children)	2,052	1,963	1,943	572 (132 schools)
Educational attainment	--	Preschool	Grade 2	Grade 5

Source: Young Lives Study (Peru).

All the information was merged into a single dataset at the child level. Minimum information requirements for inclusion were: (i) having cognitive skill scores for rounds 2 and 3; and (ii) attending a school included in the school survey.³⁸ This produced a sample of 487 children distributed in 124 schools.

Consistent with the production function described in section 2.3, period 1 variables will correspond to influences relevant from birth up to age 5, and period 2 variables will

³⁸ The risk of selection bias due to this second condition is very small. Primary school attendance in Peru is close to 100% (only 0.7% of Young Lives younger cohort children were not attending school in Round 3). In addition, schools participating in the school survey were randomly selected within the four strata considered by the designers of the survey (urban-private, urban-public, rural-public, rural-bilingual-public; see Guerrero et al. (2012)). Even so, the second requirement for inclusion will be relaxed in some of the specifications that follow in order to explore whether the results are affected by restricting the sample to those children whose school was included in the school survey.

correspond to influences relevant between the ages of 5 and 8. Accordingly, period 1 variables will be provided by rounds 1 and 2, while period 2 variables will be provided by round 3. Influences captured in the school survey (collected two years after round 3) will be assumed to be the same as those present in period 2 (i.e. I am assuming that school characteristics have not changed significantly during the two year period that separates round 3 from the school survey). I am also assuming that the child has remained in the same school since her enrolment in Grade 1 (at age 6) until the school survey was conducted (when she was 10 years old).³⁹

Rounds 2 and 3 of the Peruvian Young Lives Younger Cohort child survey contain several measures of cognitive achievement: the Cognitive Development Assessment (CDA), the Peabody Picture Vocabulary Test (PPVT), the Early Grade Reading Assessment (EGRA), a maths test, and single items in reading, writing and multiplication. This analysis will focus on the results obtained in the PPVT. This is a widely used test of receptive vocabulary that has a strong positive correlation with several measures of intelligence (Cueto and Leon, 2012). The test has a Spanish version adapted for Latin America (Dunn et al., 1986) and is the only cognitive skill measure for which the Younger Cohort survey includes longitudinal results.

Rounds 1, 2 and 3 of the household and child survey contain rich information on household and caregiver characteristics. Information related to the child is also fairly comprehensive, including aspects related to her health care and health status, schooling history, and time use (round 3). Table 2.5 presents the variables from the child and household surveys considered as early childhood and home inputs, health inputs, predetermined direct influences, and input determinants.

³⁹ According to administrative data collected from the schools included in the survey, school switching is not significant. On average, only 2% of students enrolled in primary education changed school each year between 2009 and 2010.

Regarding early childhood and home inputs, the information on household expenditure allows one to approximate the flows invested in the child in rounds 2 and 3. Since these inputs refer to those with an educational nature, the expenditure flows considered are those that account for learning materials and entertainment. For period 1, it is possible to approximate expenditure flows invested in the child and to account for access to preschool by age 5 (round 2). The survey, however, is not very informative of the quality of care and the home environment during infancy (round 1). Mothers' access to antenatal care (where advice on parenting practices is usually provided) and the way mothers' responded to the child crying were used to approximate early educational inputs or the quality of early stimulation provided to the child at home.

Period 2 (round 3) information regarding educational home inputs is richer. In addition to the approximate expenditure flows invested in the child, it is possible to account for the child's access to learning materials and resources such as books and computers, parental engagement in educational activities (i.e. providing help with homework), and the amount of time the child devoted to studying at home.

The variable chosen to reflect health inputs in both periods is an indicator of whether the child was stunted or not. This provides a fairly objective summary indicator of the child's health status. In addition, the causal relation between this measure of nutritional status and cognitive skill has already been documented (Outes-Leon et al., 2011).

Unlike the educational home inputs, the number of school characteristics that could be potentially considered within the school inputs category is notably large. In Guerrero et al. (2012), the researchers who designed and implemented the school survey distinguish between school-quality and school-responsiveness variables when describing the "educational opportunities" offered to children at school. Within the school-quality group,

variables that are potentially relevant for this analysis can be further classified into five categories: (i) size, organization and timetable (ORG); (ii) infrastructure (INF); (iii) climate (CLIM); (iv) activities and materials (ACT); and (v) teacher characteristics (TEA).

The first category includes variables capturing the number of students, teachers and classes, the presence of support staff, the number of shifts and teaching hours per day, and teacher attendance. Infrastructure comprises access to basic services and learning facilities. Variables in the school and class climate category include teachers' perception of the relations among students and between students and teachers, and of the problems and difficulties encountered during the school year. The activities and materials category comprises information regarding the use of specific learning materials in class and the degree of curriculum coverage attained. Finally, teacher characteristics include their qualifications and experience.

School responsiveness is related to the degree with which schools respond to students' needs and potential. Variables within this category (RESP) indicate whether or not the school provides support for students lagging behind or at risk of dropping out.

The school survey contains several candidate variables within each of the six categories described above. Because of this, the following three-step procedure was followed to select the variables included as school inputs: (i) pairwise correlations between candidate variables within each category were evaluated, variables with correlation coefficients below 0.6 were chosen and those with a correlation above 0.6 with two or more others were discarded; (ii) a regression of PPVT scores on the variables chosen after (i) was run for each category, and variables with a significant partial correlation were chosen; and (iii) a regression of PPVT scores on the variables chosen after (ii) was run, and those with a significant partial

correlation were chosen.⁴⁰ Variables presented in Table 2.5 within the school inputs group are the ones which resulted from this procedure.

Predetermined direct influences refer to variables that are outside the current choice set of families and that have a direct effect on skill. These include caregiver's educational attainment and age,⁴¹ and child characteristics such as sex, age and mother tongue. These child characteristics can accommodate biological differences between children as well as potential advantages that could be present when solving the test used to measure cognitive skill.⁴²

Finally, exogenous input determinants include variables reflecting family resources, parental preferences regarding the child's educational attainment, child and sibling characteristics that can affect parental investments, and an urban/rural indicator. This last variable is intended to capture differences in the general health environment and in the availability of educational goods and services in the two geographical domains.

⁴⁰ The main objective of the procedure was to reduce the amount of noise within the school inputs category as I seek a reasonably comprehensive set of school and teacher characteristics to account for the contribution of school inputs, in general, to the cognitive achievement gap. This analysis does not aim at identifying the effect of a particular school input or to rank school inputs in terms of their importance for cognitive skill formation. However, the procedure retains the dimensions of school quality considered in the school survey as this can allow one to identify particular school quality predictors that could direct further research on the individual causal effects of school-level inputs. As a robustness check, I also performed a principal component analysis for each school quality category. The results reported in the next section are robust to using the first two principal components of the six school quality dimensions (which explain between 52% and 80% of the variance). These results are presented in Appendix 2.4.

⁴¹ Caregiver engagement in educational activities will likely have a different outcome in terms of child skill formation, depending on the caregivers' parenting skills.

⁴² This means that these variables appear as relevant influences when moving from the theoretical relation which explains A_{i2} to the empirical specification for T_{i2} . For example, the PPVT is a vocabulary test, so children who do not have Spanish as their mother tongue can face a disadvantage. Differences in cognitive skill driven by some child characteristics (e.g. sex or age) can also be due to differences in unobserved inputs operating through parental preferences. This is the indirect effect of predetermined direct influences discussed above and will be captured within the contribution of the "predetermined direct influences and omitted inputs" category.

Table 2.5
Description of the variables used in the empirical specifications

Variable type	Variable used in empirical specifications	Database
Period 1 measured cognitive skill (T_{i1})	Standardized raw PPVT score	Round 2
Period 2 measured cognitive skill (T_{i2})	Standardized raw PPVT score ^(a)	Round 3
Early childhood educational inputs (HI_{i1}^0)	Real expenditure on child (learning materials and entertainment; x1,000 soles; 2006 prices in urban Lima)	Round 2
	Mother had antenatal visits during pregnancy (yes = 1)	Round 1
	Maternal response to child crying was affectionate (yes = 1) ^(b)	Round 1
	Child attended formal preschool (yes = 1)	Round 2
Period 2 educational home inputs (HI_{i2}^0)	Real expenditure on child (learning materials and entertainment; x1,000 soles; 2006 prices in urban Lima)	Round 3
	Household had books and child was encouraged to read (yes = 1)	Round 3
	Household had a computer (yes = 1)	Round 3
	Child received help from parents when doing homework (yes = 1)	Round 3
	Hours in a typical day the child spent playing	Round 3
	Hours in a typical day the child spent sleeping	Round 3
	Hours in a typical day the child spent studying	Round 3
Period 1 health input (h_{i1})	Child was stunted (yes = 1) ^(c)	Round 2
Period 2 health input (h_{i2})	Child was stunted (yes = 1)	Round 3
School inputs (SI_{i2}^0, SY_{i2})	Years of schooling (basic education)	Round 3
	Hours in a typical day the child spent at school ^(d)	Round 3
	CLIM: absence of problems in class (score 12-48) ^(e)	School survey
	INF: school had basic services (yes = 1) ^(f)	School survey
	ACT: average curricular coverage in maths and language (average % of topics covered in depth) ^(e)	School survey
	ORG: teacher absenteeism (%) ^(g)	School survey
	ORG: school had a psychologist (yes = 1)	School survey
	ORG: school was “multigrade” (yes = 1) ^(h)	School survey
	TEA: more than 50% of teachers graduated from a university (yes = 1) ^(e)	School survey
Predetermined direct influences (f_i)	Child’s caregiver has higher education (yes = 1)	Round 3
	Caregiver’s age	Round 3
	Child is male (yes = 1)	Round 3
	Child’s mother tongue is Spanish (yes = 1)	Round 3
	Child’s age in months	Round 3
Exogenous input determinants (z_i)	Child lived in urban area (yes = 1)	Round 3
	Average household total income (x10,000 soles; 2006 prices in urban Lima)	Rounds 2 and 3
	Average household size	Rounds 1, 2 and 3

Variable type	Variable used in empirical specifications	Database
	Proportion of male siblings	Rounds 1, 2 and 3
	Child birth order	Rounds 1, 2 and 3
	Caregiver aspiration for child's educational attainment was university education (yes=1)	Rounds 2 and 3

(a) Round 3 and round 2 raw PPVT scores were standardized using the round 2 mean and standard deviation.

(b) Mother cuddled or soothed child when he/she cried.

(c) A child is considered stunted if she exhibits a height for age z score below -2.

(d) The effects of children's time use categories are measured with respect to time spent working (the omitted time use category).

(e) As reported by maths and language teachers in charge of classes attended by Young Lives children.

(f) Basic services comprise water (from a public network or pipe), sanitation (public network connection or a treated cesspool), electricity and telephone connection.

(g) Measured by observation, in maths and language classes attended by Young Lives children.

(h) "Multigrade" means that children from different grades receive classes at the same time, in the same room, and from the same teacher.

Table 2.6 presents descriptive statistics as well as urban/rural differences for all the variables described above. Significant differences between urban and rural children are present in most of the direct influences and input determinants considered. Children in the urban domain enjoy larger expenditure flows from their households, have had more access to preschool and have more opportunities for learning activities at home. In addition, they attend schools with better infrastructure, more qualified teachers and greater curriculum coverage.⁴³ Finally, urban children also have better health status (are less likely to be stunted) and have more educated caregivers.

⁴³ Urban and rural averages reported in Table 2.6 should not be interpreted as characterizing the average urban and rural school. The database is at the child level so figures reported in the abovementioned table correspond to the "average school" attended by Young Lives children living in urban and rural areas. Characteristics of those schools with more Young Lives children enrolled are given more weight when characterizing this "average school".

Table 2.6
Descriptive statistics and urban-rural gaps

	Mean	SD	Urban	Rural	Diff.
Standardized raw PPVT score (round 3)	1.780	0.951	2.095	1.028	1.067*** (0.14)
Standardized raw PPVT score (round 2)	0.024	0.968	0.355	-0.766	1.121*** (0.13)
Real expenditure on child (learning materials and entertainment; round 2) ^(a)	0.274	0.364	0.342	0.112	0.23*** (0.049)
Mother had antenatal visits during pregnancy (yes = 1)	0.828	0.378	0.848	0.778	0.071* (0.038)
Maternal response to child crying was affectionate (yes = 1)	0.230	0.421	0.286	0.097	0.188*** (0.05)
Child attended formal preschool (yes = 1)	0.766	0.424	0.892	0.465	0.427*** (0.055)
Household had books and child was encouraged to read (yes = 1)	0.450	0.498	0.478	0.382	0.096 (0.06)
Household had a computer (yes = 1)	0.140	0.347	0.195	0.007	0.188*** (0.039)
Real expenditure on child (learning materials and entertainment; round 3) ^(a)	0.432	0.572	0.517	0.230	0.287*** (0.063)
Child received help from parents when doing homework (yes = 1)	0.665	0.472	0.758	0.444	0.314*** (0.029)
Hours in a typical day the child spent playing	4.346	1.517	4.488	4.005	0.483** (0.218)
Hours in a typical day the child spent sleeping	9.931	0.978	9.988	9.796	0.192 (0.114)
Hours in a typical day the child spent studying	1.945	0.834	2.120	1.526	0.594*** (0.078)
Child was stunted (yes = 1; round 2)	0.316	0.465	0.207	0.576	-0.369*** (0.034)
Child was stunted (yes = 1; round 3)	0.189	0.392	0.120	0.354	-0.235*** (0.041)
Hours in a typical day the child spent at school	6.171	0.720	6.131	6.269	-0.138 (0.108)
Years of schooling (basic education)	2.374	0.544	2.429	2.243	0.186** (0.085)
CLIM: absence of problems in class (score 12-48)	32.736	6.567	33.760	30.298	3.462** (1.317)

	Mean	SD	Urban	Rural	Diff.
INF: school had basic services (yes = 1)	0.556	0.497	0.761	0.069	0.691*** (0.087)
ACT: average curricular coverage (% of topics covered in depth)	0.531	0.153	0.564	0.452	0.111*** (0.034)
ORG: teacher absenteeism (%)	0.025	0.111	0.012	0.057	-0.045 (0.031)
ORG: school had a psychologist (yes = 1)	0.179	0.383	0.248	0.014	0.234* (0.109)
ORG: school was “multigrade” (yes = 1)	0.187	0.390	0.073	0.458	-0.385*** (0.084)
TEA: more than 50% of teachers graduated from a university (yes = 1)	0.456	0.499	0.551	0.229	0.322*** (0.091)
Child’s caregiver has higher education (yes = 1)	0.179	0.383	0.245	0.021	0.224*** (0.037)
Caregiver’s age	34.569	6.843	34.172	35.514	-1.342 (0.804)
Child is male (yes = 1)	0.478	0.500	0.490	0.451	0.038 (0.048)
Child’s mother tongue is Spanish (yes = 1)	0.893	0.309	0.985	0.674	0.312** (0.104)
Child’s age in months	96.510	3.708	96.500	96.537	-0.037 (0.507)
Child lived in urban area (yes = 1)	0.704	0.457	1.000	0.000	1.000
Average household total income ^(a)	1.512	1.116	1.711	1.037	0.674*** (0.111)
Average household size	5.538	1.849	5.270	6.176	-0.906** (0.306)
Proportion of male siblings	0.495	0.333	0.490	0.506	-0.016 (0.026)
Child birth order	2.475	1.584	2.194	3.144	-0.949*** (0.198)
Caregiver aspiration for child was university education (yes=1)	0.655	0.476	0.743	0.444	0.299*** (0.065)

The number of observations is 487 for all variables.

(a) x 1,000 soles; 2006 prices in urban Lima.

Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Results presented in Table 2.6 corroborate that the uneven distribution of direct influences affecting the skill formation process is pervasive across inputs related to different environments and originated at different periods. This, in turn, confirms that the analysis of the relative importance of different input categories when explaining unequal skill outcomes between urban and rural children is far from trivial.

2.5 Results and discussion

In this section I present and discuss the results obtained after estimating the contributions presented in Table 2.3, using the specifications and variables described in tables 2.2 and 2.5, respectively.

Table 2.7 presents normalized contributions (i.e. contributions divided by the urban-rural gap in cognitive skill at age 8) and their standard errors. Figure 2.3 shows the same point estimates accompanied by 95% confidence intervals. Panels (B) and (D) in Figure 2.3, also present the statistic and corresponding p-value of the test of omitted inputs described in the previous section. Recall that this statistic provides an estimate of the contribution of exogenous input determinants to the gap under analysis. Table 2A in Appendix 2.2 presents coefficient estimates for the variables involved in all four specifications.

Table 2.7

**Normalized contributions to the urban-rural gap in cognitive skill at age 8
(% of urban-rural gap)**


	<i>Cumulative (CU)</i>		<i>Value added plus (VAP)</i>	
<i>Prod. function</i>	Early childhood and home inputs	0.199*** (0.053)	Period 2 home inputs	0.057 (0.046)
	Health inputs	0.073*** (0.026)	Period 2 health inputs	0.038 (0.027)
	School inputs	0.540*** (0.066)	School inputs	0.370*** (0.069)
	Predetermined direct influences	0.141*** (0.032)	Period 2 predetermined direct influences	0.122*** (0.033)
	--	--	Past influences	0.400*** (0.07)
<i>Hybrid</i>	Early childhood and home inputs	0.129*** (0.049)	Period 2 home inputs	0.030 (0.041)
	Health inputs	0.063** (0.029)	Period 2 health inputs	0.041 (0.029)
	School inputs	0.479*** (0.086)	School inputs	0.348*** (0.081)
	Predetermined direct influences and omitted inputs	0.328*** (0.089)	Period 2 predetermined direct influences and period 2 omitted inputs	0.214** (0.082)
	--	--	Past influences	0.368*** (0.081)

Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Figure 2.3
Normalized
contributions to the
urban-rural gap in
cognitive skill at
age 8 (point
estimates and 95%
confidence
intervals)

Notes:
“EC&HomeInputs” refers to
early childhood and home
inputs.
“PredInf” refers to
predetermined direct influences
“PredInf&Omit” refers to
predetermined direct influences
and omitted inputs.
“P2HomeInputs” refers to
period 2 home inputs.
“P2HealthInputs” refers to
period 2 health inputs.
“PastInf” refers to past
influences.
“P2PredInf&Omit” refers to
period 2 predetermined direct
influences and period 2 omitted
inputs.



As already discussed, the hybrid-value added plus model stands out as the preferred specification based on the assumptions and error structures presented in section 2.4.1. I will, however, consider the results provided by all four specifications. This can be useful in two ways. First, it will allow one to determine whether the decompositions behave in a way that is consistent with the reasons for preferring, a priori, the hybrid-value added plus specification. These reasons refer to the presence of omitted inputs that can bias decomposition results and to the possibility of partially controlling for these unobservable influences using observable exogenous input determinants and lagged skill. Second, it can provide additional insights about the nature of omitted inputs.

Let us start with the results provided by the hybrid-cumulative specification (see panel B in Figure 2.3). School inputs stand out, with a contribution of nearly 48% of the gap, followed by the category capturing the contribution of predetermined direct influences and omitted inputs (33%). In fact, rejection of the hypothesis $(\bar{z}_U - \bar{z}_R)' \psi = 0$ indicates the presence of omitted inputs and means that exogenous input determinants included in this model have a significant contribution (22% of the gap; $p < 0.05$).

Decomposition results provided by the production function-cumulative specification (panel A in Figure 2.3) show somewhat larger contributions for the categories grouping early childhood and home inputs, and school inputs. This suggests that included school inputs can partially pick-up the contribution of the omitted influences so, if these do not belong to the school environment, the estimate provided by the production function-cumulative specification will be overstating the contribution of school influences. The hybrid-cumulative specification has the potential of mitigating this bias⁴⁴ but it retains the full

⁴⁴ Something that depends on the degree of correlation between non-observable components of input demand functions.

(cumulative) effect of unobserved innate ability, another potential source of positive bias in the estimated contribution of inputs.

Let us now turn to the results provided by the value added plus models. As already discussed, these specifications can control for period 1 omitted inputs and partially control for innate ability.⁴⁵ The first thing worth noticing is that in the hybrid- value added plus specification it is no longer possible to reject the hypothesis that exogenous input determinants have no significant contribution (see Panel D in Figure 2.3).

Failure to reject the null $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$ does not directly imply the absence of omitted inputs. However, we can combine this result with the fact that: (i) the lack of significance of exogenous input determinants in the hybrid-value added plus model is not because predetermined direct influences are a sufficient statistic to characterize the demand equation of omitted inputs⁴⁶ as these input determinants have a significant contribution in the hybrid-cumulative model; and (ii) the inclusion of lagged cognitive test scores causes an important reduction in the contribution of school inputs (compare panels A and C, and B and D in Figure 2.3) by reducing the estimated effect of several school influences (compare columns 1 and 2, and 3 and 4 in Table 2A in Appendix 2.2).

Put together, all this evidence suggests: (i) that there are period 1 omitted inputs that can be controlled for using lagged test scores; (ii) that lagged test scores can control for assignment mechanisms that correlate with other unobservables such as innate ability (as already documented in several studies reviewed in Singh (2015)); and (iii) that there are no significant period 2 omitted inputs (i.e. $\tau_1 = 0$ in equations (2.11) to (2.14)).

⁴⁵ It is worth noting that significant coefficient estimates for period 1 influences in the valued added plus specifications (see columns (2) and (4) in Table 2A in Appendix 2.2) reject the restriction that the effect of all period 1 influences decay at the same rate ρ .

⁴⁶ If predetermined direct influences are a sufficient statistic to characterize the demand function of omitted inputs, it is possible to accept the null $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$ despite the presence of omitted inputs.

It is also worth noting that these results are consistent with the characteristics of the available data, which comprise rich information on period 2 home and school influences but lack direct information on the quality of the early home environment.

Let us now focus on the results obtained after excluding the school quality inputs contained in the school survey. This implies that the only school inputs considered in the estimation and decomposition are years of schooling and hours per day spent at school. This exercise offers an additional opportunity to review the consistency of our results with respect to the assumptions that allow us to identify the contributions of interest, as well as the possibility to check their robustness to consideration of the complete sample of children who have PPVT scores in rounds 2 and 3. Results are presented in Figures 2.4 and 2.5 and coefficient estimates in tables 2B and 2C in Appendix 2.2.

If the contribution of omitted inputs is captured by their demand equations when exogenous input determinants are included in the regression, omission of significant school inputs should increase the contribution of the “predetermined direct influences and omitted inputs” category in the two hybrid specifications. Consistent with this, the contribution of this category in the hybrid-cumulative specification grows twice as large when school survey information is omitted (from 33% in the original decomposition up to 65% according to panel B in Figure 2.4). There is also strong evidence of the presence of omitted inputs as the contribution of exogenous input determinants in the hybrid cumulative model is now 51% (it was 22% in the original decomposition) and highly significant ($p < 0.00$).

A value added plus specification allows a reduction in the contribution of the category hosting omitted inputs (down to 42%; see panel (D) in Figure 2.4) and this is consistent with the presence of omitted period 1 influences, something that has already been suggested by the preceding analysis. Importantly, however, the contribution of exogenous input

determinants in the hybrid-value added plus decomposition remains significant (30%; $p < 0.00$) which implies we cannot accept the null $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$. This result, which differs from the one obtained with the complete set of data, confirms there are still relevant period 2 influences omitted. This is consistent with the fact that we are intentionally omitting school inputs.

It is also worth noting that the production function-cumulative specification provides now a much larger estimate of the contribution of early childhood and home inputs than when considering the full set of variables (40% vs. 20%), and that the school inputs considered in these set of results only account for around 7% of the gap (see panel A in Figure 2.4). This means that the school inputs considered are failing to pick-up the contribution of their omitted counterparts. This is consistent with the fact that years of schooling and hours per day spent at school are fairly similar across children (between 2 and 3 years and around 6 hours, respectively) and, thus, are weakly correlated with the characteristics of schools and teachers we are omitting.⁴⁷

Results presented in Figure 2.5 reveal that the results just discussed are robust to considering the entire sample of children with complete PPVT scores and not just those attending schools included in the school survey. This should mitigate concerns regarding potential selection bias in the sample used for the main analysis. The use of a larger sample also adds precision to the results discussed in the previous three paragraphs.

Finally, in Appendix 2.3 I also present decomposition results from a hybrid-value added plus model after instrumenting lagged PPVT test scores with the results of the Cognitive Developmental Assessment (CDA) test, also administered to the younger cohort during the second round (see Cueto and Leon (2012) for details). The objective is to remove the

⁴⁷ This is just another expression of the fact that, in Peru, relevant inequalities in education are in terms of quality and not in terms of access.

attenuation bias due to measurement error from our estimates of the persistence parameter, and to provide an upper bound for the contribution of past influences and a lower bound for the contribution of school inputs. As expected, the instrumental variable estimate of the persistence parameter is larger than its OLS counterpart (see Table 2D in Appendix 2.3). Accordingly, the contribution of past influences is now situated around 50% while school inputs contribute 28% of the gap (see Table 2E in Appendix 2.3).

Based on all the above, and focusing on the results provided by the hybrid- value added plus models, we can confidently conclude that around 35 per cent (and no less than 28 per cent) of the gap in cognitive skill between urban and rural 8 year old children in Peru is related to differences in school inputs (years of schooling, school and teacher characteristics) received between ages 6 and 8. Differences in past influences, mostly related to the learning and care environment to which the child was exposed during early childhood, contribute no more than 50% to this gap.

Figure 2.4
Normalized
contributions to the
urban-rural gap in
cognitive skill at
age 8 - excluding
school quality
inputs provided by
the school survey

Notes:

“EC&HomeInputs” refers to early childhood and home inputs.

“PredInf” refers to predetermined direct influences

“PredInf&Omit” refers to predetermined direct influences and omitted inputs.

“P2HomeInputs” refers to period 2 home inputs.

“P2HealthInputs” refers to period 2 health inputs.

“PastInf” refers to past influences.

“P2PredInf&Omit” refers to period 2 predetermined direct influences and period 2 omitted inputs.





Figure 2.5
Normalized
contributions to
the urban-rural
gap in cognitive
skill at age 8 -
excluding school
inputs from the
school survey and
working with the
complete sample

Notes:
 “EC&HomeInputs” refers to early childhood and home inputs.
 “PredInf” refers to predetermined direct influences
 “PredInf&Omit” refers to predetermined direct influences and omitted inputs.
 “P2HomeInputs” refers to period 2 home inputs.
 “P2HealthInputs” refers to period 2 health inputs.
 “PastInf” refers to past influences.
 “P2PredInf&Omit” refers to period 2 predetermined direct influences and period 2 omitted inputs.



2.6 Concluding remarks

This analysis aimed at measuring the contribution of school and early childhood influences to the difference in cognitive development observed, at age 8, between urban and rural children in Peru. The study has contributed new evidence based on an unusually rich dataset and provided by a decomposition strategy less prone to biases than those used so far in the literature. Two key features of the strategy that allow it to be less susceptible to biased decomposition results are the explicit distinction between skill inputs and skill input determinants, and the use of a special variable category to host omitted influences.

As discussed in the preceding analysis, including skill input determinants to control for omitted influences can be useful when decomposing a gap, but only as long as we are aware of the assumptions we are making when we assign these controls to a particular category. If family resources and preferences, and local supply characteristics determine the quantity and quality of inputs received both at home and at school, these controls will capture the contribution of omitted inputs that belong to both environments so a bias can arise if we assign them exclusively to one of them. As will be further explored in the next chapter, previous analysis are susceptible to this type of bias as they have typically assigned the contribution of all variables reflecting household and family characteristics to the category hosting home influences.

The results obtained in this analysis indicate that observable school influences occurring between ages 6 and 8, account for a significant share of urban/rural cognitive skill gap (around 35% and no less than 28%). The share attributable to differences in these children's early environment (between ages 0 and 5) is no larger than 50%.

These results have implications for the interpretation of available evidence regarding the evolution of cognitive achievement gaps. Recall that, in Peru, as in several other developing countries, these gaps emerge before children enter school and remain practically unchanged during their school years. This evidence has been interpreted as indicating that early childhood influences play the leading role in shaping these gaps, so little can be achieved by investing after children have reached school age (Heckman, 2006; Schady et al., 2014). In terms of interpretation, our results show that the fact that cognitive achievement gaps were already present at preschool age does not imply that schools play a minor role. The reason for this is that skill exhibits less than perfect persistence so only a fraction of the previous gap remains every period. In fact, the value added models provide a persistence parameter estimate between 0.34 and 0.52. This means that between 70 and 80 per cent of skill persists each year.⁴⁸

Judging that schools play a minor role in shaping achievement gaps based on the fact that these gaps were already present and had a similar size before children entered school can be interpreted as using a difference-in-difference approach to estimate the effect of being “treated” with a particular type of school (in this case, a rural school). In terms of a value added model, this approach implies assuming perfect persistence in skill; that is, it implies introducing the restriction $\rho = 1$ and modelling the first difference. In Andrabi et al. (2011), the authors discuss how imposing this restriction when it is not true will lead to a downward bias in the estimated effect of contemporaneous inputs.⁴⁹ This is precisely what happens when skill exhibits less than perfect persistence (as in our case) but one compares the size of achievement gaps at school age against the size of gaps present before children began

⁴⁸ Considering that 3 years have passed between the two measurements of skill and that $0.7^3 = 0.34$.


⁴⁹ The reader can also consult Angrist and Pischke (2009) (Appendix 5.4) on the difference between assuming a standard fixed effect model (e.g. $T_{i2} = x'_{i2}\beta_1 + \alpha_i + \varepsilon_{i2}$) and an autoregressive process (e.g. $T_{i2} = \rho T_{i1} + x'_{i2}\beta_1 + \varepsilon_{i2}$) and the biases that will arise when estimating treatment effects (β_1) by first differencing when the true process is the latter. In our case, the autoregressive process is consistent with the cumulative nature of skill formation.

school: one gets the erroneous impression that schools do not matter. As opposed to this, the results presented in this chapter show that observable differences between rural and urban schools⁵⁰ do matter and the fact that the gap in cognitive skill was already present at age 5 implies that these differences are serving to sustain the inequalities in cognitive development that emerged before children were enrolled at school.

⁵⁰ The urban-rural indicator used in this analysis corresponds to the domain where the child lives. This means that the relevant difference is between schools attended by urban and rural children. Drawing a conclusion in terms of urban and rural schools, however, is acceptable as 95% of rural children in the sample attend a rural school, while 99% of urban children in the sample attend an urban school.


Figure 2.1

Proportion of second grade students with a satisfactory attainment in reading comprehension (urban and rural point estimates and 95% confidence intervals)



Source: Control Sample of the National Student Evaluation (2007-2014), Ministry of Education - Peru.

Figure 2.2
Urban-rural differences in cognitive achievement
(standardized Peabody Picture Vocabulary Test scores)



	Round 2	Round 3	Obs.
Urban	0.29	2.00	1,276
Rural	-0.74	0.96	493
Gap	1.03***	1.04***	

Significant at the 1% (***), 5% (**), 10% (*).

Source: Own calculations based on the Peruvian Young Lives database (younger cohort; rounds 2 and 3).

Appendix 2.1

Consistent estimation of the contribution of an input category

To ease manipulation, let us assume that there are only two relevant categories of direct influences:⁵¹ early childhood and school inputs. Further assume there is only one early childhood influence (HI_{i1}) and two school inputs (SI_{i2}^O, SI_{i2}^U), one of which remains unobserved (SI_{i2}^U). As usual, I assume innate ability (μ_{i0}) is unobserved and skill is measured with random error (ε_{i2}), although I will ignore their potential correlation with observed inputs in order to focus on the biases caused by the omission of one of the school inputs. The empirical version of the production function of skill is given by:

$$T_{i2} = HI_{i1}\gamma_2 + SI_{i2}^O\phi_1^O + e_i \quad (i)$$

where $e_i = SI_{i2}^U\phi_1^U + \mu_{i0}\beta_{(2)} + \varepsilon_{i2}$.

Under this setting, the contributions of early childhood and school inputs to the urban-rural gap in cognitive skill observed in period 2 are given by:

$$CE_2 = (E(HI_{i1}|U) - E(HI_{i1}|R))\gamma_2 \quad (ii)$$

$$CS_2 = (E(SI_{i2}^O|U) - E(SI_{i2}^O|R))\phi_1^O + (E(SI_{i2}^U|U) - E(SI_{i2}^U|R))\phi_1^U \quad (iii)$$

where expectations are taken conditioned on the child belonging to the urban ($E(\cdot|U)$) or rural ($E(\cdot|R)$) domain.

The estimation of production function parameters given in (i) will be affected by biases caused by the omission of SI_{i2}^U if this school input is correlated with the other direct influences. Let us express the relationship between the omitted input and the rest of direct influences by means of the following linear projection:

$$SI_{i2}^U = b_1HI_{i1} + b_2SI_{i2}^O + w_i \quad (iv)$$

An OLS regression of (iv) would allow us to recover the partial correlation coefficients: $\hat{b}_1$ and $\hat{b}_2$. These coefficients determine the sign and size of the biases affecting the OLS estimates of the parameters given in (i). In particular:⁵²

$$plim \hat{\gamma}_2 = \gamma_2 + \phi_1^U plim(\hat{b}_1) \quad (v)$$

$$plim \hat{\phi}_1^O = \phi_1^O + \phi_1^U plim(\hat{b}_2) \quad (vi)$$

⁵¹ This analysis can be extended to the more general case considered in the main text without affecting its main results.

⁵² The expressions that follow are just the probability limits of a special case of the following more general (and familiar) result. Assume $y = X_1\beta_1 + X_2\beta_2 + \varepsilon$. In a regression that omits X_2 , the OLS estimate of β_1 is given by: $\hat{\beta}_1 = (X_1'X_1)^{-1}X_1'y = (X_1'X_1)^{-1}X_1'(X_1\beta_1 + X_2\beta_2 + \varepsilon) = \beta_1 + (X_1'X_1)^{-1}X_1'X_2\beta_2 + (X_1'X_1)^{-1}X_1'\varepsilon$. The $k_1 \times k_2$ matrix $(X_1'X_1)^{-1}X_1'X_2$ contains the OLS coefficients of a regression of each variable in X_2 on all variables in X_1 . In our case, $X_1 = [HI_{i1} \ SI_{i2}^O]$ and $X_2 = [SI_{i2}^U]$. Therefore, $\beta_1 = (\gamma_2 \ \phi_1^O)'$ and $(X_1'X_1)^{-1}X_1'X_2\beta_2 = (\hat{b}_1 \ \hat{b}_2)'\phi_1^U$. The assumption of absence of correlation between included inputs and innate ability and measurement error would correspond to the result $plim (X_1'X_1)^{-1}X_1'\varepsilon = 0$ in this general case.

From expressions (v) and (vi) is clear that there must be no correlation between the omitted input and included influences ($b_1 = b_2 = 0$) for OLS to yield consistent estimates of their direct effects.

A situation where $b_1 = 0$ and $b_2 \neq 0$ implies that the omitted input exhibits partial correlation only with the observed influence that belongs to its same category (i.e. the school environment). In other words, the included school input is able to fully pick-up the effect of its omitted counterpart. Under these circumstances we will not be able to obtain a consistent estimate of the effect of the observed school influence (SI_{i2}^U will act as a confounder) but we will still obtain a consistent estimate of the contribution of the school inputs category. To see this, consider that an OLS regression of (i) would allow one to estimate the contribution of school inputs using:

$$\widehat{CS}_2 = (\overline{SI}_{U2}^O - \overline{SI}_{R2}^O) \hat{\phi}_1^O \quad (\text{vii})$$

Where, for a given variable X_{it} , $\overline{X}_{Ut}$ and $\overline{X}_{Rt}$ correspond to the sample counterparts of $E(X_{it}|U)$ and $E(X_{it}|R)$, respectively. Based on (vi), this estimate exhibits the following probability limit:

$$\begin{aligned} \text{plim } \widehat{CS}_2 &= (E(SI_{i2}^O|U) - E(SI_{i2}^O|R)) (\phi_1^O + \phi_1^U \text{plim}(\hat{b}_2)) \\ &= (E(SI_{i2}^O|U) - E(SI_{i2}^O|R)) \phi_1^O + (E(SI_{i2}^O|U) - E(SI_{i2}^O|R)) \phi_1^U b_2 \end{aligned} \quad (\text{viii})$$

Following (iv) and considering that $b_1 = 0$ and $b_2 \neq 0$, we can express the urban/rural difference in the average unobserved school input as follows:

$$E(SI_{i2}^U|U) - E(SI_{i2}^U|R) = b_2 (E(SI_{i2}^O|U) - E(SI_{i2}^O|R)) \quad (\text{ix})$$

Finally, combining (viii) and (ix) yields:

$$\text{plim } \widehat{CS}_2 = (E(SI_{i2}^O|U) - E(SI_{i2}^O|R)) \phi_1^O + (E(SI_{i2}^U|U) - E(SI_{i2}^U|R)) \phi_1^U = CS_2 \quad (\text{x})$$

The expression given in (x) indicates that we can obtain a consistent estimate of the contribution of school inputs despite an omitted input is acting as a confounder for the estimation of individual effects within this category.

Appendix 2.2

Coefficient estimates for the four empirical specifications

Table 2A

Coefficient estimates for the four specifications including school inputs from the school survey

VARIABLES	(1) Prod.Fn-CU	(2) Prod.Fn-VAP	(3) Hybrid-CU	(4) Hybrid-VAP
Real expenditure in child (learning materials and entertainment; round 2)	0.125 (0.095)	0.021 (0.086)	0.076 (0.090)	-0.022 (0.083)
Mother had antenatal visits during pregnancy (yes = 1)	0.137** (0.055)	0.129** (0.052)	0.132** (0.051)	0.121** (0.052)
Maternal response to child cry was affectionate (yes = 1)	0.086 (0.064)	0.030 (0.070)	0.077 (0.063)	0.025 (0.070)
Child attended formal preschool (yes = 1)	0.088 (0.083)	0.020 (0.071)	0.049 (0.094)	0.003 (0.080)
Household has books and child is encouraged to read (yes = 1)	0.220*** (0.063)	0.242*** (0.069)	0.210*** (0.057)	0.235*** (0.068)
Household has a computer (yes = 1)	0.106* (0.059)	0.074 (0.054)	0.087 (0.059)	0.064 (0.055)
Real expenditure in child (learning materials and entertainment; round 3)	0.063 (0.068)	0.035 (0.069)	0.062 (0.062)	0.036 (0.064)
Child receives help from parents when doing homework (yes = 1)	0.087 (0.098)	0.022 (0.097)	0.007 (0.093)	-0.041 (0.091)
Hours in a typical day the child spends playing	0.004 (0.026)	-0.009 (0.029)	-0.003 (0.021)	-0.012 (0.026)
Hours in a typical day the child spends sleeping	-0.028 (0.038)	-0.044 (0.038)	-0.032 (0.036)	-0.044 (0.038)
Hours in a typical day the child spends studying	0.062 (0.046)	0.033 (0.040)	0.045 (0.045)	0.024 (0.041)
Child is stunted (yes = 1; round 2)	-0.083 (0.075)	-0.026 (0.078)	-0.048 (0.085)	-0.002 (0.084)
Child is stunted (yes = 1; round 3)	-0.200* (0.112)	-0.173 (0.123)	-0.212 (0.123)	-0.184 (0.133)
Hours in a typical day the child spends at school	-0.075 (0.049)	-0.067 (0.051)	-0.075 (0.051)	-0.070 (0.053)
Years of schooling (basic education)	0.348*** (0.082)	0.252*** (0.077)	0.342*** (0.081)	0.253*** (0.075)
CLIM: absence of problems in class (score 12-48)	0.011** (0.004)	0.012** (0.004)	0.009** (0.004)	0.011** (0.004)
INF: school has basic services (yes = 1)	0.229*** (0.064)	0.053 (0.065)	0.183** (0.069)	0.044 (0.064)
ACT: average curricular coverage (% of topics covered in depth)	0.612* (0.308)	0.432 (0.261)	0.521 (0.308)	0.388 (0.267)
ORG: teacher absenteeism (%)	-0.862* (0.416)	-0.830** (0.323)	-0.780* (0.380)	-0.761** (0.317)
ORG: school has a psychologist (yes = 1)	0.173* (0.085)	0.187* (0.090)	0.194** (0.079)	0.203** (0.087)
ORG: school is “multigrade” (yes = 1)	-0.320** (0.119)	-0.326*** (0.091)	-0.292** (0.120)	-0.308*** (0.098)

VARIABLES	(1) Prod.Fn-CU	(2) Prod.Fn-VAP	(3) Hybrid-CU	(4) Hybrid-VAP
TEA: more than 50% of teachers graduated from university (yes = 1)	0.102** (0.047)	0.016 (0.047)	0.090* (0.049)	0.012 (0.046)
Child's caregiver has higher education (yes = 1)	0.187** (0.064)	0.055 (0.049)	0.135* (0.066)	0.017 (0.056)
Caregiver's age	0.000 (0.004)	-0.002 (0.004)	0.013** (0.005)	0.008 (0.005)
Child is male (yes = 1)	0.013 (0.039)	-0.018 (0.035)	0.028 (0.096)	0.034 (0.083)
Child's mother tongue is Spanish (yes = 1)	0.348*** (0.113)	0.370*** (0.120)	0.341** (0.144)	0.389** (0.152)
Child's age in months	0.013 (0.012)	0.003 (0.009)	0.014 (0.011)	0.004 (0.009)
Child lives in urban area (yes = 1)			0.120 (0.113)	0.028 (0.096)
Average household total income			0.018 (0.024)	0.010 (0.021)
Average household size			0.013 (0.026)	0.013 (0.026)
Proportion of male siblings			-0.032 (0.175)	-0.100 (0.173)
Child birth order			-0.097*** (0.026)	-0.085** (0.031)
Caregiver aspiration for child is university education (yes = 1)			0.058 (0.054)	0.026 (0.059)
Standardized raw PPVT score (round 2)		0.348*** (0.049)		0.341*** (0.052)
Constant	-1.259 (1.249)	0.572 (1.145)	-1.362 (1.176)	0.465 (1.117)
Observations	487	487	487	487
R-squared	0.547	0.601	0.557	0.608

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 2B

Coefficient estimates for the four specifications excluding school inputs from the school survey

VARIABLES	(1) Prod.Fn-CU	(2) Prod.Fn-VAP	(3) Hybrid-CU	(4) Hybrid-VAP
Real expenditure in child (learning materials and entertainment; round 2)	0.281** (0.124)	0.099 (0.093)	0.180 (0.106)	0.038 (0.087)
Mother had antenatal visits during pregnancy (yes = 1)	0.143** (0.054)	0.129** (0.050)	0.136** (0.050)	0.119** (0.053)
Maternal response to child cry was affectionate (yes = 1)	0.132* (0.071)	0.047 (0.071)	0.101 (0.062)	0.035 (0.065)
Child attended formal preschool (yes = 1)	0.184 (0.107)	0.078 (0.098)	0.050 (0.121)	0.009 (0.113)
Household has books and child is encouraged to read (yes = 1)	0.229*** (0.063)	0.258*** (0.068)	0.207*** (0.052)	0.240*** (0.062)
Household has a computer (yes = 1)	0.195** (0.079)	0.129* (0.072)	0.133* (0.075)	0.099 (0.072)
Real expenditure in child (learning materials and entertainment; round 3)	0.088 (0.069)	0.041 (0.070)	0.077 (0.058)	0.041 (0.064)
Child receives help from parents when doing homework (yes = 1)	0.141 (0.107)	0.058 (0.101)	0.027 (0.085)	-0.023 (0.085)
Hours in a typical day the child spends playing	0.066** (0.023)	0.030 (0.028)	0.032* (0.017)	0.015 (0.022)
Hours in a typical day the child spends sleeping	0.023 (0.039)	-0.012 (0.037)	-0.011 (0.032)	-0.028 (0.035)
Hours in a typical day the child spends studying	0.154*** (0.043)	0.081* (0.042)	0.085* (0.046)	0.047 (0.046)
Child is stunted (yes = 1; round 2)	-0.179* (0.087)	-0.082 (0.088)	-0.079 (0.101)	-0.025 (0.094)
Child is stunted (yes = 1; round 3)	-0.219* (0.118)	-0.190 (0.132)	-0.226 (0.134)	-0.198 (0.143)
Hours in a typical day the child spends at school	-0.049 (0.060)	-0.045 (0.054)	-0.042 (0.056)	-0.043 (0.054)
Years of schooling (basic education)	0.342*** (0.083)	0.225*** (0.063)	0.326*** (0.077)	0.227*** (0.060)
Child's caregiver has higher education (yes = 1)	0.260*** (0.081)	0.075 (0.060)	0.169** (0.075)	0.025 (0.058)
Caregiver's age	-0.003 (0.004)	-0.005 (0.004)	0.012** (0.005)	0.007 (0.005)
Child is male (yes = 1)	-0.005 (0.044)	-0.037 (0.036)	0.056 (0.094)	0.062 (0.078)
Child's mother tongue is Spanish (yes = 1)	0.505*** (0.115)	0.479*** (0.131)	0.390** (0.145)	0.439** (0.170)
Child's age in months	0.016 (0.013)	0.005 (0.010)	0.018 (0.011)	0.007 (0.010)
Child lives in urban area (yes = 1)			0.395*** (0.076)	0.202** (0.091)
Average household total income			0.033 (0.032)	0.021 (0.027)
Average household size			0.006 (0.028)	0.006 (0.028)

VARIABLES	(1) Prod.Fn-CU	(2) Prod.Fn-VAP	(3) Hybrid-CU	(4) Hybrid-VAP
Proportion of male siblings			-0.137 (0.146)	-0.202 (0.138)
Child birth order			-0.101*** (0.025)	-0.088*** (0.029)
Caregiver aspiration for child is university education (yes = 1)			0.112 (0.072)	0.065 (0.070)
Standardized raw PPVT score (round 2)		0.399*** (0.038)		0.363*** (0.046)
Constant	-2.008 (1.347)	0.228 (1.255)	-1.894 (1.152)	0.102 (1.172)
Observations	487	487	487	487
R-squared	0.474	0.556	0.510	0.573

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 2C

Coefficient estimates for the four specifications excluding school inputs from the school survey and working with the complete sample

VARIABLES	(1) Prod.Fn-CU	(2) Prod.Fn-VAP	(3) Hybrid-CU	(4) Hybrid-VAP
Real expenditure in child (learning materials and entertainment; round 2)	0.170*** (0.043)	0.050 (0.030)	0.122*** (0.036)	0.034 (0.028)
Mother had antenatal visits during pregnancy (yes = 1)	0.188*** (0.048)	0.134*** (0.038)	0.193*** (0.044)	0.139*** (0.039)
Maternal response to child cry was affectionate (yes = 1)	0.055 (0.054)	0.015 (0.051)	0.042 (0.043)	0.012 (0.046)
Child attended formal preschool (yes = 1)	0.158** (0.059)	0.062 (0.045)	0.059 (0.058)	0.016 (0.044)
Household has books and child is encouraged to read (yes = 1)	0.219*** (0.059)	0.216*** (0.053)	0.183*** (0.048)	0.194*** (0.048)
Household has a computer (yes = 1)	0.226*** (0.041)	0.126*** (0.036)	0.130*** (0.040)	0.082** (0.034)
Real expenditure in child (learning materials and entertainment; round 3)	0.081* (0.045)	0.037 (0.030)	0.059 (0.044)	0.029 (0.032)
Child receives help from parents when doing homework (yes = 1)	0.165*** (0.044)	0.111** (0.046)	0.076 (0.045)	0.061 (0.044)
Hours in a typical day the child spends playing	0.067*** (0.021)	0.035** (0.016)	0.037* (0.021)	0.021 (0.018)
Hours in a typical day the child spends sleeping	0.009 (0.025)	-0.024 (0.023)	-0.019 (0.024)	-0.036 (0.024)
Hours in a typical day the child spends studying	0.130*** (0.024)	0.044** (0.018)	0.065*** (0.022)	0.017 (0.018)
Child is stunted (yes = 1; round 2)	-0.154** (0.070)	-0.044 (0.050)	-0.098 (0.066)	-0.024 (0.050)
Child is stunted (yes = 1; round 3)	-0.164** (0.060)	-0.131** (0.057)	-0.125** (0.052)	-0.111** (0.053)
Hours in a typical day the child spends at school	0.036 (0.035)	0.011 (0.033)	0.020 (0.037)	0.005 (0.034)
Years of schooling (basic education)	0.294*** (0.045)	0.142*** (0.035)	0.254*** (0.044)	0.134*** (0.034)
Child's caregiver has higher education (yes = 1)	0.308*** (0.061)	0.103* (0.052)	0.194*** (0.058)	0.054 (0.048)
Caregiver's age	-0.002 (0.003)	-0.001 (0.003)	0.009** (0.004)	0.006 (0.004)
Child is male (yes = 1)	0.061 (0.036)	0.042 (0.029)	0.068 (0.059)	0.060 (0.038)
Child's mother tongue is Spanish (yes = 1)	0.513*** (0.076)	0.495*** (0.064)	0.278** (0.105)	0.367*** (0.100)
Child's age in months	0.010 (0.008)	0.007 (0.006)	0.012 (0.007)	0.008 (0.006)
Child lives in urban area (yes = 1)			0.453*** (0.085)	0.246*** (0.083)
Average household total income			0.048** (0.018)	0.027* (0.014)
Average household size			0.009 (0.014)	0.008 (0.013)

VARIABLES	(1) Prod.Fn-CU	(2) Prod.Fn-VAP	(3) Hybrid-CU	(4) Hybrid-VAP
Proportion of male siblings			-0.012 (0.069)	-0.029 (0.052)
Child birth order			-0.076*** (0.020)	-0.051** (0.021)
Caregiver aspiration for child is university education (yes = 1)			0.216*** (0.034)	0.147*** (0.033)
Standardized raw PPVT score (round 2)		0.447*** (0.027)		0.399*** (0.030)
Constant	-1.848** (0.753)	-0.135 (0.750)	-1.667* (0.807)	-0.249 (0.813)
Observations	1,561	1,561	1,561	1,561
R-squared	0.425	0.542	0.473	0.557

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Appendix 2.3

Estimation and decomposition results after instrumenting lagged PPVT scores with CDA test scores

Table 2D

Coefficient estimates for the hybrid-value added plus model (OLS and IV)

VARIABLES	Hybrid-VAP OLS	Hybrid-VAP IV ^(a)
Real expenditure in child (learning materials and entertainment; round 2)	-0.022 (0.083)	-0.073 (0.098)
Mother had antenatal visits during pregnancy (yes = 1)	0.121** (0.052)	0.115** (0.052)
Maternal response to child cry was affectionate (yes = 1)	0.025 (0.070)	-0.002 (0.071)
Child attended formal preschool (yes = 1)	0.003 (0.080)	-0.020 (0.071)
Household has books and child is encouraged to read (yes = 1)	0.235*** (0.068)	0.248*** (0.070)
Household has a computer (yes = 1)	0.064 (0.055)	0.053 (0.056)
Real expenditure in child (learning materials and entertainment; round 3)	0.036 (0.064)	0.023 (0.060)
Child receives help from parents when doing homework (yes = 1)	-0.041 (0.091)	-0.065 (0.091)
Hours in a typical day the child spends playing	-0.012 (0.026)	-0.016 (0.028)
Hours in a typical day the child spends sleeping	-0.044 (0.038)	-0.050 (0.040)
Hours in a typical day the child spends studying	0.024 (0.041)	0.013 (0.039)
Child is stunted (yes = 1; round 2)	-0.002 (0.084)	0.022 (0.093)
Child is stunted (yes = 1; round 3)	-0.184 (0.133)	-0.169 (0.130)
Hours in a typical day the child spends at school	-0.070 (0.053)	-0.067 (0.050)
Years of schooling (basic education)	0.253*** (0.075)	0.207*** (0.081)
CLIM: absence of problems in class (score 12-48)	0.011** (0.004)	0.012** (0.005)
INF: school has basic services (yes = 1)	0.044 (0.064)	-0.028 (0.099)
ACT: average curricular coverage (% of topics covered in depth)	0.388 (0.267)	0.319 (0.203)
ORG: teacher absenteeism (%)	-0.761** (0.317)	-0.751*** (0.272)
ORG: school has a psychologist (yes = 1)	0.203** (0.087)	0.207** (0.088)
ORG: school is “multigrade” (yes = 1)	-0.308*** (0.098)	-0.316*** (0.084)

VARIABLES	Hybrid-VAP OLS	Hybrid-VAP IV ^(a)
TEA: more than 50% of teachers graduated from university (yes = 1)	0.012 (0.046)	-0.029 (0.065)
Child's caregiver has higher education (yes = 1)	0.017 (0.056)	-0.043 (0.080)
Caregiver's age	0.008 (0.005)	0.006 (0.006)
Child is male (yes = 1)	0.034 (0.083)	0.037 (0.075)
Child's mother tongue is Spanish (yes = 1)	0.389** (0.152)	0.415*** (0.152)
Child's age in months	0.004 (0.009)	-0.002 (0.007)
Child lives in urban area (yes = 1)	0.028 (0.096)	-0.020 (0.092)
Average household total income	0.010 (0.021)	0.006 (0.021)
Average household size	0.013 (0.026)	0.013 (0.025)
Proportion of male siblings	-0.100 (0.173)	-0.135 (0.175)
Child birth order	-0.085** (0.031)	-0.078** (0.033)
Caregiver aspiration for child is university education (yes = 1)	0.026 (0.059)	0.010 (0.066)
Standardized raw PPVT score (round 2)	0.341*** (0.052)	0.517*** (0.165)
Constant	0.465 (1.117)	1.406 (1.159)
Observations	487	487
R-squared	0.608	0.594

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

(a) IV estimation uses round 2 Cognitive Developmental Assessment (CDA) test scores as instrument for round 2 PPVT test scores. CDA coefficient in first stage = 0.268 (p<0.01).

Table 2E**Decomposition results for the hybrid-value added plus model (OLS and IV)**

	<i>Value added plus - OLS</i>		<i>Value added plus – IV (a)</i>	
Hybrid	Period 2 home inputs	0.030 (0.041)	Period 2 home inputs	0.009 (0.048)
	Period 2 health inputs	0.041 (0.029)	Period 2 health inputs	0.037 (0.030)
	School inputs	0.348*** (0.081)	School inputs	0.280*** (0.091)
	Period 2 predetermined direct influences and period 2 omitted inputs	0.214** (0.082)	Period 2 predetermined direct influences and period 2 omitted inputs	0.154 (0.095)
	Past influences	0.368*** (0.081)	Past influences	0.519*** (0.155)


(a) IV estimation uses round 2 Cognitive Developmental Assessment (CDA) test scores as instrument for round 2 PPVT test scores. CDA coefficient in first stage = 0.268 (p<0.01).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix 2.4


Main decomposition results using the first two principal components of each school quality dimension

(A) Hybrid - value added plus (original school inputs considered in the main analysis)



Ho: $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$; stat = 0.11, p-value = 0.266

(B) Hybrid - value added plus (first two principal components of the six school quality dimensions)



Ho: $(\bar{z}_U - \bar{z}_R)' \tilde{\psi} = 0$; stat = 0.09, p-value = 0.352

Notes:

“EC&HomeInputs” refers to early childhood and home inputs.

“PredInf” refers to predetermined direct influences

“PredInf&Omit” refers to predetermined direct influences and omitted inputs.

“P2HomeInputs” refers to period 2 home inputs.

“P2HealthInputs” refers to period 2 health inputs.

“PastInf” refers to past influences.

“P2PredInf&Omit” refers to period 2 predetermined direct influences and period 2 omitted inputs.

Chapter 3

Linear Decompositions of Cognitive Achievement Gaps: A Cautionary Note and an Illustration Using Peruvian Data

Abstract

Developmental gaps between children of different socioeconomic backgrounds emerge early and persist over time. Linear decompositions based on the Oaxaca-Blinder (OB) technique are a fairly common way of estimating the contribution of two or more categories of variables to these differences in cognitive achievement. Two prominent examples of these categories are family and school influences. In this regard, the literature shows no consensus in terms of how to implement these OB decompositions and interpret their components. It also exhibits a tendency to separate home and school influences by assigning all observed household, family and child characteristics to the first category. I argue that this can lead to misleading policy implications and to biases in the estimated contributions of the categories. This analysis seeks to contribute to the literature in two ways. First, it formally explores the potential for biases in the decomposition exercises attempted thus far. Second, it offers an alternative decomposition strategy consistent with explicit behavioural assumptions regarding the determination of skill inputs. This prevents arbitrary choices in terms of empirical strategy, number of components and their interpretation, and also makes the analysis less prone to biases. I illustrate empirically the main points of the analysis employing a rich dataset that contains longitudinal information on cognitive test scores, family and school characteristics, to decompose the cognitive skill gap observed, at age 8, between urban and rural children in Peru.

3.1 Introduction and motivation

Developmental gaps between children from disadvantaged backgrounds and those belonging to more affluent families emerge early and persist over time (Heckman, 2006, 2007; Paxson and Schady, 2007b; Schady et al., 2014; Walker et al., 2007). Evidence suggests that such differences are difficult to overcome later in life, and limit these children's future economic opportunities and wellbeing (Almond and Currie, 2011; Cunha et al., 2006).

Cognitive skill formation is a cumulative process and, thus, all relevant influences that had taken place before skill is measured can, in principle, play role in shaping these gaps. A relevant question that follows concerns which particular influence or group of influences play a significant role for the emergence of these differences. Do earlier influences matter more than those occurring later in the life of these children? Do influences originating in a particular environment (such as these children's home or school) play a major part?

The literature has tried to address this type of questions in several ways. One way has been to estimate the individual effects of particular influences, reporting their size and significance. This is the strand of literature focusing on the "technology of skill formation". The main empirical challenge related to the estimation of individual effects on skill is related to the presence of unobserved influences. Omitted influences are likely to generate biased estimates of individual effects because skill inputs are choice variables and are interrelated through the decision making process of families. This strand of the literature has explicitly exposed this problem by laying down models that postulate a production function of skill and characterize how families' choices determine its inputs (Cunha and Heckman, 2007; Glewwe and Miguel, 2008; Todd and Wolpin, 2003).

Identification of individual effects in this literature has usually relied on some form of instrumental variable strategy. This is feasible because of the limited number of parameters of interest. As already documented in Chapter 1, findings using the US National Longitudinal Survey of the Young (NLSY/79) confirm that skill formation is a cumulative process, that socio-emotional skills affect cognitive skill, and that cognitive skill is particularly sensitive to parental investments during early childhood (Cunha and Heckman, 2008; Cunha, Heckman, and Schenach, 2010; Todd and Wolpin, 2007). Efforts to replicate this in the developing world have confirmed the importance of parental investments and the

fact that cognitive and socio-emotional skills are related (Helmerts and Patnam, 2011; Lopez-Boo, 2009).

This strand of the literature has also made important contributions by making explicit the assumptions required by different empirical specifications to identify production function parameters (Todd and Wolpin, 2003, 2007) and by clarifying the difference between the partial and the total effect of an input on skill (Glewwe and Miguel, 2008; Todd and Wolpin, 2003).⁵³

Another strand of the literature has attempted a more direct answer to the question of which influence or group of influences is more important for the emergence of a particular gap, by proposing a linear decomposition of this gap. In particular, the difference in mean outcomes between two groups of children is decomposed into contributions that are due to the differences in the mean values of two or more categories or groups of variables. These categories are usually built in order to compare the relative importance of influences originated at home and those originated at school (Hernandez-Zavala et al., 2006; McEwan, 2004; McEwan and Marshall, 2004; McEwan and Trowbridge, 2007; Ramos et al., 2012).

Most of the studies in this strand of the literature have relied on some form of Oaxaca-Blinder decomposition technique (Blinder, 1973; Oaxaca, 1973). There are several ways to implement this technique and this implies that one needs to choose which specific strategy to follow and how to interpret its components. In addition, one has to devise a rule to classify variables and contributions into different categories. In this regard, a review of the studies applied to the developing world reveals two problematic features: (i) that the choice

⁵³ The partial or marginal effect of an input corresponds to its production function parameter. It implies one is holding all other direct influences constant. The total effect of an input corresponds to its partial effect plus those that occur through the changes in other inputs caused by the shift in the input of interest. Under the logic a model describing families' choices, the total effect corresponds to the parameter in a conditional demand function (see Glewwe and Miguel (2008)). As discussed in Chapter 2, experimental and quasi-experimental methods usually recover the total effect of an input.

of decomposition strategy and interpretation of its components has been made arbitrarily (i.e. with no indication of the assumptions in terms of skill formation process and family behaviour that led to these choices); and (ii) that the rule commonly employed to separate home and school influences has been to assign all observed household, family and child characteristics to the first category.

The source of these two problematic features is the lack of a decomposition strategy based on the predictions of a framework describing the production of skill and the process determining its inputs (this is, the lack of a decomposition strategy that takes into account the insights and lessons of the “technology of skill formation” literature). The problem with these two features is that they entail the risk of producing misleading policy implications and of introducing biases in the estimated contributions of the categories of interest. This potential source of bias has been overlooked so far in the literature and emerges because several of the home and family characteristics considered within the “home influences” category can control for omitted inputs that belong to the group of “school influences”.

A good example of the above is family income or wealth. Under the logic of a model describing the production of skill and families’ choices, family income has no direct effect on skill but acts as an input determinant. In fact, family income cannot only determine the quantity and quality of inputs received at home but also the quantity and quality of inputs received at school. If the latter is true, it would not be appropriate to attribute the contribution of this variable exclusively to the “family influences” category.

Based on the above, this analysis seeks to contribute to the literature by formally exploring the two problematic features of the decomposition exercises attempted so far, and by offering an alternative decomposition strategy consistent with explicit behavioural assumptions regarding the determination of skill inputs. The latter will prevent arbitrary

choices in terms of decomposition technique, its components and interpretation, and will make the analysis less susceptible to biases. This chapter will, therefore, focus on the framework behind the empirical strategy employed in Chapter 2. I will analyse how this strategy compares to the Oaxaca-Blinder technique and will discuss in which ways it is superior to the decomposition exercises found in the literature.

The rest of the chapter is organized as follows. Section 3.2 reviews the methods and recent results associated with the decomposition exercises attempted so far for the developing world. Section 3.3, presents a framework describing the skill formation process and how families' choices determine its inputs, allowing for endogenous school quality. In section 3.4, I use the insights of this model to formally explore the potential biases that can be introduced by the decomposition strategies employed in the literature. I also use these insights to propose an alternative decomposition strategy less prone to these biases and to discuss its rationale under the lens of the Oaxaca-Blinder technique. In section 3.5, I use a rich dataset comprising cognitive test scores, family background and school information for 8-year-old Peruvian children to decompose the urban/rural gap in cognitive development, and empirically illustrate the main points made in section 3.4. Section 3.5 shares some results with the previous chapter but complements and compares them against the results of a simulation exercise that follows the decomposition rules commonly employed in the literature. Section 3.6 closes with some final remarks.

3.2 Decomposing achievement gaps in developing countries: arbitrary choices and potential biases

Table 3.1 (at the end of this chapter) summarizes a comprehensive list of studies that have attempted a linear decomposition of the differences in average cognitive achievement between two groups of children living in the developing world. Differences in average cognitive outcomes are expressed as a linear combination of differences in the averages of

predictors, and the contributions of different subsets of predictors are estimated. These predictors or influences are typically grouped into categories that comprise family and school characteristics (see column D in Table 3.1).

The results of these exercises are summarized in column F of Table 3.1. A striking feature of these results is their lack of robustness regarding the contribution of school characteristics to the difference in cognitive outcomes between children of different backgrounds. In McEwan and Trowbridge (2007) and McEwan (2004), for example, the authors analysed learning outcome gaps between indigenous and non-indigenous children in Guatemala, Bolivia and Chile. They concluded that differences in the quality of schools make a significant contribution to these gaps, explaining between 50 and 70%. In Ramos et al. (2012), the authors analysed the difference in PISA results between urban and rural students in Colombia. They also found that differences in the school environment play a significant role with a contribution that ranges between 75 and 83% of the observed gap.

Results presented in Hernandez-Zavala et al. (2006) tell quite a different story. These authors addressed learning outcome gaps between indigenous and non-indigenous children in Guatemala, Mexico and Peru. They concluded that differences in “family variables” contribute more than differences in “school variables” to the overall explained gap. Surprisingly, they found that the contribution of school characteristics in Guatemala ranges between 17 and 23%, which is in sharp contrast with the results discussed in McEwan and Trowbridge (2007) where the contribution of schools to the same gap was found to be as high as 70%.

In a similar fashion, and although they do not consider an explicit category containing school variables,⁵⁴ in Arteaga and Glewwe (2014), the authors highlight the role played by household and child characteristics above that of community characteristics. They analysed cognitive test score gaps between indigenous and non-indigenous children in Peru and found that, by age 8 (when children are in Grade 2), differences in household and child characteristics account for 80% of the gap.

Further analysis of the studies reviewed in Table 3.1 reveals two additional characteristics that can help explain the lack of consensus regarding the importance for cognitive achievement gaps of differences in school characteristics *vis-à-vis* the importance of differences in family variables. First, the choice of decomposition strategy and the interpretation of its components (summarized in column C of Table 3.1) can be characterized as arbitrary. In other words, there is no explicit reference to the assumptions that have led to choosing a particular empirical strategy to decompose the observed gap and to interpret the components into which it has been decomposed.

Second, in all the studies providing an estimate of the contribution of different subsets of observed influences (e.g. family and school characteristics), control variables have been assigned to particular subsets or categories (as reported in columns D and E) without consideration of the role they play in the production of skill. The potentially harmful consequences of these two features are: (i) the risk of introducing biases in the estimate of the contribution of particular categories of variables; and (ii) the risk of overlooking the role of relevant influences when discussing policy implications. In what follows, I further develop these ideas.

⁵⁴ The authors, however, present their community level fixed effects as partially capturing differences in school and teacher characteristics.

3.2.1 Choosing a decomposition strategy and interpreting its components

All the studies summarized in Table 3.1 have used some variant of the Oaxaca-Blinder, henceforth OB, decomposition. In general, this strategy is based on decomposing the difference in mean outcomes between two groups into a portion due to differences in the mean values of observable predictors, and a portion due to differences in the coefficients governing the relationship between the outcome and these predictors. The latter is usually described as the “unexplained” part of the gap.

There are different ways to implement the OB decomposition. Depending on the number of components involved in the decomposition, these are usually classified as “twofold” or “threefold” (Jann, 2008). Let us start by exploring the “threefold” decompositions. For this, consider two groups of individuals (A and B) for whom a certain outcome (y_i) can be related to a set of predictors (x_i) in the following way:

$$\begin{aligned}y_{iA} &= x'_{iA}\beta_A + \varepsilon_{iA} \\y_{iB} &= x'_{iB}\beta_B + \varepsilon_{iB}\end{aligned}\tag{3.1}$$

If we estimate a linear regression for each group including an intercept in both x_{iA} and x_{iB} , the following will hold: $\bar{y}_A = \bar{x}'_A\hat{\beta}_A$ and $\bar{y}_B = \bar{x}'_B\hat{\beta}_B$.

One way of measuring how much of the difference in mean outcomes has to do with differences in predictors and how much with differences in coefficients is by using the following decomposition (Jones and Kelly, 1984; Winsborough and Dickenson, 1971):

$$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_B + \bar{x}'_B(\hat{\beta}_A - \hat{\beta}_B) + (\bar{x}_A - \bar{x}_B)'(\hat{\beta}_A - \hat{\beta}_B)\tag{3.2}$$

This is a “threefold” decomposition that takes group B as the reference group. The first component captures the contribution of the difference in predictors or endowments (the portion of the gap that would be closed if group B had the same endowments as group A). The second component captures the contribution of the difference in coefficients (the portion of the gap that would be closed if group B had the same coefficients as group A). The third component is an interaction term that accounts for the fact that differences in endowments and coefficients occur simultaneously. It is the portion of the gap that only arises if endowments and returns change together (Biewen, 2012).

If the decomposition takes group A as the reference group, it yields:

$$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_A + \bar{x}_A'(\hat{\beta}_A - \hat{\beta}_B) - (\bar{x}_A - \bar{x}_B)'(\hat{\beta}_A - \hat{\beta}_B) \quad (3.3)$$

The interpretation is similar to that provided in the previous paragraph but with changes occurring in the endowments and coefficients of group A.

Let us now briefly focus on the “twofold” decompositions. These are better appreciated if we introduce a third vector of reference coefficients ($\hat{\beta}_R$) to measure the contributions of the differences in endowments and coefficients to the overall gap. The difference in mean outcomes between groups A and B can be expressed as follows:

$$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_R + \bar{x}_A'(\hat{\beta}_A - \hat{\beta}_R) + \bar{x}_B'(\hat{\beta}_R - \hat{\beta}_B) \quad (3.4)$$

The four types of “twofold” decompositions usually encountered in the literature emerge depending on the choice of $\hat{\beta}_R$ (see Table 3.2).

Table 3.2
Four different two-fold Oaxaca-Blinder decompositions

	Reference coefficients	Decomposition
1	Group A coefficients ($\hat{\beta}_R = \hat{\beta}_A$)	$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_A + \bar{x}_B' (\hat{\beta}_A - \hat{\beta}_B)$
2	Group B coefficients ($\hat{\beta}_R = \hat{\beta}_B$)	$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_B + \bar{x}_A' (\hat{\beta}_A - \hat{\beta}_B)$
3	Coefficients from a pooled regression over both groups ($\hat{\beta}_R = \hat{\beta}_{Pooled}$)	$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_{Pooled} + \bar{x}_A' (\hat{\beta}_A - \hat{\beta}_{Pooled}) + \bar{x}_B' (\hat{\beta}_{Pooled} - \hat{\beta}_B)$
4	Coefficients from a pooled regression over both groups including a group indicator ($\hat{\beta}_R = \hat{\beta}_{Pooled*}$)	$\bar{y}_A - \bar{y}_B = (\bar{x}_A - \bar{x}_B)' \hat{\beta}_{Pooled*} + \bar{x}_A' (\hat{\beta}_A - \hat{\beta}_{Pooled*}) + \bar{x}_B' (\hat{\beta}_{Pooled*} - \hat{\beta}_B)$

In all four cases, the first term in the right hand side of the decomposition equation captures the portion of the gap that can be explained by the differences in endowments. The remaining term(s) capture the portion of the gap that is due to differences in coefficients or the “unexplained” part of the gap.

The first two “twofold” decompositions correspond to the original formulations proposed in Oaxaca (1973) and Blinder (1973). The third decomposition was considered in Neumark (1988). This author analysed the implications in terms of firm behaviour of choosing different reference groups when estimating the contribution of “discrimination” to a wage gap. He proposed using the coefficients of a pooled regression as a reference as this requires weaker assumptions in terms of firms’ discriminatory behaviour.

Finally, the fourth “twofold” decomposition is also based on a vector of reference coefficients obtained from a pooled regression but including a group indicator. It is worth noticing that this is equivalent to estimating the “unexplained” part of the gap by using the coefficient of the group indicator in the pooled regression. To see this, consider this pooled* regression to be as follows:

$$y_i = x_i' \beta_{Pooled*} + \delta D_i + \varepsilon_i \quad (3.5)$$

where D_i is the group indicator (a dummy variable that takes the value of 1 if the individual belongs to group A and the value of 0 if he belongs to group B). The inclusion of this group indicator ensures that the regression line passes through the means of both groups. Therefore: $\bar{y}_A = \bar{x}_A' \hat{\beta}_{Pooled*} + \hat{\delta}$ and $\bar{y}_B = \bar{x}_B' \hat{\beta}_{Pooled*}$. This, in turn, implies that:

$$\begin{aligned} \bar{y}_A - \bar{y}_B &= (\bar{x}_A - \bar{x}_B)' \hat{\beta}_{Pooled*} + \hat{\delta} \\ &= (\bar{x}_A - \bar{x}_B)' \hat{\beta}_{Pooled*} + \bar{x}_A' (\hat{\beta}_A - \hat{\beta}_{Pooled*}) + \bar{x}_B' (\hat{\beta}_{Pooled*} - \hat{\beta}_B) \end{aligned} \quad (3.6)$$

which means that $\bar{x}_A' (\hat{\beta}_A - \hat{\beta}_{Pooled*}) + \bar{x}_B' (\hat{\beta}_{Pooled*} - \hat{\beta}_B) = \hat{\delta}$.

There is no consensus regarding which is the best OB decomposition. In Biewen (2012), for example, the author advocates for the “threefold” strategy arguing that the interaction term is a constituent part of the difference in means and “it is hard to find reasons to allocate [it] either in whole or in part to either the ‘characteristics’ or the ‘returns’ effect” (Biewen, 2012; p. 12). Its interpretation, however, can be quite problematic, especially if it accounts for a substantial portion the overall gap.

Within the family of “twofold” decompositions, in Elder et al. (2010) and Jann (2008) the authors advocate for the OB pooled* option. In both studies, the authors argue that the alternative OB pooled option tends to overstate the contribution of the “explained” component. In Elder et al. (2010) the authors further show that the “unexplained” part estimated using the OB pooled* decomposition is usually close to the estimates provided by the more standard OB options (those where the reference coefficients are set to the coefficients of group A or B). They, therefore, propose the OB pooled* strategy as an attractive method for obtaining single measures of the “explained” and “unexplained” portions of a gap.

This lack of consensus is manifest in the literature surveyed in Table 3.1. Out of the ten studies surveyed, two use the “threefold” decomposition (studies 6 and 10); two use a “twofold” decomposition taking the disadvantaged group as reference (3 and 8); one uses a “twofold” decomposition taking the advantaged group as reference (5); one uses a “twofold” decomposition and reports results taking the advantaged and disadvantaged groups as reference (4); three use a “twofold” decomposition using the coefficients of a pooled regression as reference (1,2,8), and two of them use the pooled* option (1 and 2); and one study uses both a “twofold” decomposition taking the advantaged group as reference and one taking the coefficients of a pooled* regression as reference (study number 7).

There are also different interpretations given to the size and significance of the “unexplained” part of the gap. For example, some studies explicitly acknowledge that the “unexplained” part of the gap can be capturing the contribution of omitted influences (Hernandez-Zavala et al., 2006; McEwan, 2004; McEwan and Marshall, 2004; McEwan and Trowbridge, 2007). Other studies, however, implicitly assume that the skill formation process has been fully specified, and interpret the difference in coefficients capturing the “unexplained” part of the gap as literally revealing a difference in the effectiveness with which inputs are transformed into skill by the two groups under analysis (Barrera-Osorio et al., 2011; Beltran and Seinfeld, 2012; Burger, 2011; Zhang and Lee, 2011).

A common feature of all the studies presented in Table 3.1 is that the choice of decomposition strategy is not based on a framework that describes the production of skill and the process determining its direct influences or inputs. This leads to arbitrary choices in terms of the number of components and reference groups used to build these components, as well as to arbitrary interpretations of the results regarding their contribution to the gap under

analysis. This explains the lack of consensus regarding which decomposition strategy to use and can lead to misleading policy implications.

Consider, for example, those studies that interpret the “unexplained” part of the gap as revealing a difference in the effectiveness with which inputs are transformed into skill. Without an explicit reference to a production function of skill, and in case the “unexplained” part dominates, this interpretation will lead to policy recommendations that advocate for a more efficient use of school resources instead of an increase in the provision of school inputs (see, for example, Beltran and Seinfeld (2012) and Burger (2011)).⁵⁵ This type of policy recommendation can be misleading if the difference in coefficients is in fact a symptom of omitted inputs that are unevenly distributed between the two groups. In this regard, the risk of omitted inputs is particularly significant in the studies surveyed in Table 3.1, as they all rely on cross sectional data. This means that they can account only for contemporaneous influences whereas the skill formation process is cumulative, which implies that outcomes observed in a particular moment of time are a function of all influences that have occurred until that moment.

3.2.2 Building up categories and assigning contributions

In eight out of the ten studies summarized in Table 3.1, the authors go beyond the “difference in endowments-difference in coefficients” dichotomy and further decompose the former into different subsets or categories of variables. These categories typically comprise family and school characteristics. The objective, thus, is to measure the contribution of influences related to these two environments to the gap under analysis. The discussion that

⁵⁵ It is also worth noticing that this interpretation is consistent with the notion that the learning process is an attribute of the school and not an attribute of the child. Conceptualizing the skill formation technology as an attribute of the school implies that the learning process would cease in absence of the school, just as production would stop in absence of the firm. This implication is especially problematic when modelling broad forms of skill whose acquisition is a process that started before and continues beyond the schooling period.

follows is centred around empirical exercises that focus on these two categories. The main messages, however, can be generalized to situations that involve more than two categories.

Estimating the contribution of family and school influences to the gap under analysis requires an estimate of the effects of these influences on skill and a rule to assign these influences into the categories proposed. Both elements entail the risk of introducing a bias in the estimate of the contributions of family and school influences. The first source of bias has been widely addressed in the literature focused on the “technology of skill formation” and is related to the presence of omitted variable biases in the estimates of production function parameters. Because we seldom observe all the relevant direct influences of skill and these are interrelated through the decision making process of families, it is highly likely that omitted influences will produce biased estimates of the direct effects of those influences we observe. Different empirical specifications of the production function of skill require different assumptions to be able to recover these parameters. These assumptions have been discussed at length in Todd and Wolpin (2003) and Todd and Wolpin (2007).⁵⁶

As discussed in Chapter 2, the second source of bias has not been addressed yet in the literature and is related to the use of rules (explicit or implicit) that end up assigning the contribution of variables that belong to one category to another. The problem arises when assigning the contribution of variables that control (directly or indirectly) for omitted influences that belong to more than one category. If this is the case, the contribution of these controls should not be assigned exclusively to either the family or school environment. If omitted influences have a positive effect on skill and the gap in their endowment is also

⁵⁶ It is worth recalling that consistent estimation of the contribution of a subset of variables does not require consistent estimation of the effects of all observed influences. As discussed in the previous chapter, omission of a relevant input of skill can still allow one to recover a consistent estimate of the contribution of a category of variables as long as the effect of the omitted input is picked up by observed influences that belong to its same category.

positive, doing so will lead to overstating the contribution of the category hosting these controls.

Studies surveyed in Table 3.1 exhibit the abovementioned problem in two different ways. The first has to do with the assignment of predetermined household, family and child characteristics to the “family influences” category. The second is related to the use of school fixed effects or school-level averages of child characteristics and their assignment to the “school influences” category.

The inclusion of predetermined child, household and family characteristics that do not have a direct effect on skill (such as family income, household size or the child’s birth order) in the estimation of a production function is justified insofar they are relevant arguments in the demand functions of omitted direct influences. Given constraints and preferences, parents play a major role deciding the inputs that determine the skill formation process of their children. Because of this, arguments in the demand function of inputs are related to child, family and household characteristics. Inclusion of these predetermined controls implies we are replacing the omitted influences by their corresponding demand functions. This configures what is known as a “hybrid” specification (Rosenzweig and Schultz, 1983; Todd and Wolpin, 2007).

As shown in Table 3.1, a rule commonly employed in the literature when assigning variables into categories has been to group all family, household and child characteristics into the “family influences” category. A quick revision of the variables typically considered within this category reveals that these include direct influences (such as books or time that parents spend with children) but also variables that reflect family resources and preferences which are, therefore, controlling for omitted influences.

Family resources and preferences determine the inputs provided in the home environment (such as early stimulation opportunities or learning material) but can also play a role determining the quantity and quality of inputs provided at school through parents' school choices. More affluent families can provide better stimulation opportunities to their children during early childhood and can also afford enrolling them in better schools. Because of this, the rule employed in the literature entails the risk of overstating the relative importance influences provided in the family environment. This risk grows larger as families' school choices have a greater influence on the quality of school inputs and as less information on school inputs is available for the analysis.

This type of bias is likely affecting the results discussed in McEwan and Marshall (2004) and Hernandez-Zavala et al. (2006), where the authors found that family variables contribute more than schools to the gap under analysis, after grouping all family, household and child characteristics into a single category. In addition, it is worth noting that in Hernandez-Zavala et al. (2006), school data were especially limited for Mexico and the contribution of school influences was found to be zero (see Table 3.1, row 5, columns E and F).

Not all the studies that have grouped family and household characteristics into a single category have found that schools have a limited contribution. In fact, three studies that found that differences in school characteristics play a major role also followed this rule (McEwan, 2004; McEwan and Trowbridge, 2007; Ramos et al., 2012) (see rows 1, 2 and 3 in Table 3.1). A common feature of these studies that can explain this result is that they have captured school influences by introducing school fixed effects or the mean socioeconomic level of the peer group.

In principle, one could argue that the contribution of school fixed effects or school-level averages of child characteristics belong to the “school characteristics” category. School fixed effects absorb all direct influences that are invariant within schools and school inputs are surely among these. However, influences originated at school might not be the only inputs shared by students that belong to the same school. The stronger the correlation between children’s socioeconomic status and the quality of schooling received, the closer the match between children’s early childhood skill and school choice. Under this setting, poor information on early childhood inputs or past skill measures (as in the three studies mentioned above) will lead to school fixed effects or school-level averages of child characteristics absorbing omitted non-school influences, and to an overestimation of the contribution of the school environment.⁵⁷

In a similar fashion to the arbitrary choices of decomposition strategy and interpretation of its components, the problems discussed in this section can also be traced back to the lack of an explicit framework describing the skill formation process and families’ choices determining its inputs. In particular, the lack of such framework leads one to overlook the difference between skill inputs and skill input determinants (those variables that belong to the demand function of inputs). This, in turn, increases the risk of using rules that end up assigning the contribution of one category to another.

3.3 The production function of skill and families’ choices regarding its inputs

In this section I describe the skill formation technology and present a simple model describing how families’ choices determine its inputs. The objective is to formalise the

⁵⁷ Interestingly, in Ramos et al. (2012) the authors included the mean socioeconomic level of the peer group among the “school characteristics” category but in their conclusions interpreted its contribution as a family influence. This is reasonable insofar these school averages capture children’s early childhood skill, but entails the same risk of understating the importance of schools as the rule of assigning all family and household characteristics into a single category.

difference between the inputs to skill formation and the variables that determine these inputs, postulate how are they related, and describe the potential roles that input determinants can play in an empirical model seeking to explain the skill formation process. This will serve to illustrate the risk of bias if one assigns contributions to variable categories following the rule commonly employed in the literature. It will also serve to guide the design of an alternative decomposition strategy that mitigates this risk.

The production function of skill is the same as that postulated in Chapter 2. I divide the relevant phase of child development into two time periods. The first begins when the child is born and finishes at age 5. The second period corresponds to the time when the child remains within primary school age. Skill exhibited by child i at the end of period 2 (A_{i2}) is a function of contemporaneous and past direct influences affecting the child. Formally:

$$A_{i2} = A_2(HI_{i2}, HI_{i1}, SI_{i2}, SY_{i2}, h_{i2}, h_{i1}, f_i, \mu_{i0}) \quad (3.7)$$

where HI_{i1} are educational inputs provided during early childhood (period 1); HI_{i2} are educational inputs provided at home during period 2; SI_{i2} are educational inputs provided at the school where the child is enrolled during period 2; SY_{i2} are years of schooling attained during period 2; h_{it} indicates the child's health status during period t ; f_i captures predetermined direct influences; and μ_{i0} is the child's innate ability.

It is worth recalling that expression (3.7) denotes a structural relationship between skill and those variables that have a direct effect on it. The arguments in this production function are similar to those proposed in Glewwe and Miguel (2008), except for the presence of f_i . This formulation, therefore, allows for predetermined child and parental characteristics (e.g. parental education) to have a direct influence on skill. In addition, I will further classify direct influences as inputs (if they are determined by families' choices during the period under analysis) or as predetermined (if they are outside the current choice set of families).

It is worth mentioning that some studies have found a positive relationship between socio-emotional and cognitive skills working with both developed and developing country data (Cunha et al., 2006; Outes-Leon et al., 2010). This does not mean, however, that the socio-emotional development of the child is a missing argument in equation (3.7). In fact, I am assuming that socio-emotional skills are affected by the same or by a subset of the direct influences considered in (3.7). Thus, if one substitutes the socio-emotional skill vector by its corresponding influences, one would obtain the production function given in (3.7). Todd and Wolpin (2007) make a similar argument when discussing the value-added specification for cognitive skills. It implies that the parameters identified in (3.7) capture the *total* effect of the corresponding influence on cognitive skill formation.

The fact that inputs are choice variables and we seldom observe all relevant influences complicates the estimation of their effects due to endogeneity problems.⁵⁸ However, this same fact can provide important insights regarding the different types of relations than can be postulated between children's skill and its determinants. A clear understanding of these relations will play an important role in the design of a decomposition strategy that minimizes the risk of incurring in biases when building up categories and assigning contributions. For this, we first need to consider a model describing families' choices.

The model presented here follows Glewwe and Miguel (2008) closely but extends their original formulation to allow for endogenous school inputs. Glewwe and Miguel (2008) assume that school and teacher characteristics available to the child are not influenced by parental decisions made during the period under analysis (between the child's conception and the end of the primary school cycle). During this period, families' choices related to the school environment are limited to the number of years of schooling. This is consistent, for

⁵⁸ In this case, "endogeneity" is understood as the presence of correlation between observed and unobserved influences.

example, with a situation where school inputs are solely a function of the family's location decision and this decision was made prior to the period under analysis and cannot be changed. It should be noticed that if the location decision can fully characterize the school inputs available to the child, the supply of educational services within each locality must be fairly homogeneous.

In this regard, it is reasonable to assume that parents can influence the school and teacher characteristics available to their children either by changing location (migrating) or because localities are better characterized by a distribution of educational services from where parents can choose, rather than by a homogeneous type of school. Under this setting, the simplest assumption is that all families can chose a school from a common pool or choice set (see, for example, Todd and Wolpin (2003)). This is consistent with a situation where there is a similar distribution of schooling services across localities or migration costs are not significant.

In what follows, I will adopt a more flexible approach. I will assume that families can choose a particular school (j) with a particular set of characteristics (SI_{ij}) from a given set $S_i = \{SI_{i1}, \dots, SI_{ij_i}\}$. This set is not necessarily the same for all families. In general terms, the set S_i can be defined by the distribution of educational services available in the geographical area within which each family's location decisions occur, a characteristic which is specific to the context under analysis.

At this point is worth noticing that the objective of this model is not to explain how location decisions are taken and how these affect the quality of school services available to the child. The objective is to illustrate the relation between the inputs of skill and its determinants, allowing family choices to affect school characteristics considering the fact that the supply of educational services available to each family is not necessarily the same.

Consistent with the two-period setting assumed above, consider that parents maximise the following utility function:

$$U_i = U(C_{i1}, C_{i2}, h_{i2}, h_{i1}, A_{i2}; \tau, \sigma, \omega) \quad (3.8)$$

where C_{it} is child i 's parental consumption of an aggregate good in period t , and τ, σ and ω reflect parental preferences regarding time, child's skill and child's health, respectively.

Child health is determined according to the following production functions:

$$h_{i1} = H_1(c_{i1}, M_{i1}, HE_{i1}, \eta_{i0}) \quad (3.9)$$

$$h_{i2} = H_2(h_{i1}, c_{i2}, M_{i2}, HE_{i2}, \eta_{i0}) \quad (3.10)$$

where c_{it} is child i 's consumption of the aggregate good in period t , M_{it} are health inputs provided in period t , HE_{it} captures the local health environment in period t , and η_{i0} is the child's innate healthiness.

Under this setting, parents choose consumption levels (C_{it} and c_{it}), health inputs (M_{it}), educational inputs provided during early childhood and at home (HI_{i1}, HI_{i2}), and years of schooling in a particular school (SY_{i2}) to maximize utility given in (3.8), subject to the skill formation technology given in (3.7), the production functions for health given in (3.9) and (3.10), and the following budget constraint:

$$Y_{i1} - S_{i1} = p_{c1}(C_{i1} + c_{i1}) + p_{m1}M_{i1} + p_{h1}HI_{i1} \quad (3.11)$$

$$Y_{i2} + (1 + r)S_{i1} = p_{c2}(C_{i2} + c_{i2}) + p_{m2}M_{i2} + p_{h2}HI_{i2} + \sum_{j=1}^{J_i} p_{ij}^s SY_{ij} 1(SY_{ij} = SY_{i2}) \quad (3.12)$$

In (3.11) and (3.12), S_{i1} represent savings, p_{ct} is the price of the aggregate consumption good in period t , p_{mt} is the price of health inputs in period t , p_{h1} is the price of educational inputs provided during early childhood, p_{h2} is the price of educational inputs provided at home during period 2, p_{ij}^s is the price of one year of schooling at school with characteristics

SI_{ij} , SY_{ij} is the number of years of schooling demanded at school j , and $1(SY_{ij} = SY_{i2})$ is an indicator function that equals 1 in case school j has been chosen ($SY_{ij} = SY_{i2}$) and 0 otherwise.⁵⁹ Finally, Y_{it} is period t exogenously determined income, and r is the interest rate at which parents are assumed can borrow or lend between the two time periods.

As already explained, I assume that parents can choose a particular school (j) from a given set ($S_i; |S_i| = J_i$). This feature of the model implies that parents will be able not only to choose the number of years of schooling, but that they can also influence the educational inputs provided at school. In fact, by choosing a certain number of years of schooling at a particular school, parents are also determining that their child will be exposed to a certain quality of educational inputs. This means that we need an additional expression to fully characterize the optimization problem faced by parents. Formally:

$$SI_{i2} = \sum_{j=1}^{J_i} SI_{ij} 1(SY_{ij} = SY_{i2}) \quad (3.13)$$

The first order conditions of the problem stated above provide the relationships explaining the optimal levels of consumption, health inputs, educational home inputs, years of schooling and school inputs. All of these demand functions depend on: (i) resources (Y_{i1}, Y_{i2}); (ii) prices (r, p, p_{ij}^s) $j = 1, \dots, J_i$ and $p = (p_{c1}, p_{c2}, p_{m1}, p_{m2}, p_{h1}, p_{h2})$; (iii) exogenous environmental variables (HE_{i1}, HE_{i2}, S_i); (iv) predetermined direct influences (f_i); (v) endowments (μ_{i0}, η_{i0}); and (vi) preferences (τ, σ, ω). Formally:

$$C_{it}^* = C_t(Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \quad t = 1, 2; j = 1, \dots, J_i \quad (3.14)$$

$$c_{it}^* = c_t(Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \quad t = 1, 2; j = 1, \dots, J_i \quad (3.15)$$

⁵⁹ Notice that I am assuming that children do not switch schools during the period under analysis.

$$M_{it}^* = M_t(Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \\ t = 1, 2; j = 1, \dots, J_i \quad (3.16)$$

$$HI_{it}^* = H_t(Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \\ t = 1, 2; j = 1, \dots, J_i \quad (3.17)$$

$$SY_{i2}^* = SY(Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \\ j = 1, \dots, J_i \quad (3.18)$$

$$SI_{i2}^* = \sum_{j=1}^{J_i} SI_{ij} 1(SY_{ij} = SY_{i2}^*) \quad (3.19)$$

The production function indicated in (3.7) involves only, and all, of the variables that have a direct effect on skill, whether they are predetermined or not. In addition to this function, there are three other meaningful relations that can be postulated to explain children's skill: a demand function, a conditional demand function, and the hybrid production function presented in the previous chapter. In what follows I briefly describe these functions to clarify the role that exogenous input determinants can play in the estimation of the production function of skill.⁶⁰

The demand function for cognitive skill involves only predetermined variables that can have a direct or indirect effect on skill. It can be obtained by replacing (3.17), (3.18) and (3.19) in (3.7), replacing (3.15) and (3.16) in (3.9) and (3.10) and solving the demand for child's health in periods 1 and 2, and inserting these solutions into (3.7). This yields:

$$A_{i2} = A_2^D(Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \quad j = 1, \dots, J_i \quad (3.20)$$

The conditional skill demand function conditioned on input k , only involves input k and controls for the exogenous determinants of those inputs not included. To obtain this relation we need first to consider the demand functions for the rest of inputs conditioned over input k . These are obtained by fixing input k at its utility maximising level, which implies that

⁶⁰ For a more complete description of how to obtain and interpret these functions, the reader can consult Glewwe and Miguel (2008).

prices related to input k and resources devoted to its consumption are no longer relevant arguments of the demand for the rest of inputs.

For example, demand functions for educational inputs provided during early childhood and at home conditioned over school inputs (years of schooling and school characteristics) are given by:

$$HI_{it}^{CD} = H_t(SI_{i2}, SY_{i2}; Y_{CD}; p; HE_{i1}, HE_{i2}; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \quad t = 1, 2 \quad (3.21)$$

where Y_{CD} refers to resources after adjusting for school expenditures $Y_{CD} = Y_{i1} + \frac{Y_{i2}}{1+r} - \frac{\sum_{j=1}^{J_i} p_{ij}^s SY_{i2} 1(SY_{ij}=SY_{i2})}{1+r}$. As already noted, the price of schooling (p_{ij}^s) is no longer present in (3.21).

Similar expressions can be obtained for the demand for child's health in both periods after building conditional demand functions for child's consumption and health inputs. Replacing conditional demand functions for early childhood and educational home inputs and child's health in the production function given in (3.7) yields the demand for child's skill conditioned over school inputs. Formally:

$$A_{i2} = A_2^{CD}(SI_{i2}, SY_{i2}; Y_{CD}; p; HE_{i1}, HE_{i2}; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \quad (3.22)$$

Finally, and following the example centred on school inputs, a hybrid production function can be obtained if we replace all inputs in (3.7), except those related to the school environment, by their respective demand functions. Thus we obtain:

$$A_{i2} = A_2^H(SI_{i2}, SY_{i2}; Y_{i1}, Y_{i2}; r, p, p_{ij}^s; HE_{i1}, HE_{i2}, S_i; f_i; \mu_{i0}, \eta_{i0}; \tau, \sigma, \omega) \quad j = 1, \dots, J_i \quad (3.23)$$

The differences between expressions (3.7), (3.20), (3.22) and (3.23) have important consequences for empirical work. Consider, for example, the difference between the effects

of school inputs provided by equations (3.7) and (3.22). The effect of school inputs captured in equation (3.7) corresponds only to the direct impact of these inputs on skill, holding all other direct influences constant. The effect of school inputs provided by expression (3.22) includes this direct influence but also captures the indirect effect produced through changes in other inputs within the choice set of parents. Accordingly, experimental designs and instrumental variable techniques will typically identify the latter; i.e. they will typically identify the parameters of a conditional demand function (or the “policy effects” as denoted in Todd and Wolpin (2003)).⁶¹

A hybrid function such as the one given in expression (3.23) allows one to recover the parameters of the production function of observed inputs. The motivation for this type of specification is empirical and stems from the possibility of avoiding omitted variable biases originated by the presence of unobserved inputs. Under the rationale of a model of family choice such as the one described above, the use of exogenous input determinants in the estimation of a production function implies that the researcher believes in the possibility of omitted inputs and that these have been replaced by their corresponding demand functions.

3.4 The cognitive skill gap, empirical specifications and decomposition strategies

In this section I use the insights provided by the model described above to illustrate the potential biases that can be introduced by the rule commonly employed in the literature to assign the contribution of individual variables to the categories proposed for the decomposition. For this, I will consider the decomposition of the cognitive skill gap observed, in period 2, between children belonging to two generic groups (A and B). I will consider two different specifications and the feasible components related to each of them. I will start by discussing the potential biases present under the rule commonly employed in

⁶¹ Notice that in an experimental setting, after introducing exogenous variation in a certain input nothing prevents post-treatment values of other inputs to change in response to treatment.

the literature (henceforth, the “standard decomposition rule”). I will then present an alternative decomposition strategy less prone to these biases, and discuss its rationale under the lens of the Oaxaca-Blinder technique.

3.4.1 The risk of bias under the standard decomposition rule

Let us assume that the production function given in (3.7) is approximately linear. To ease the exposition, also assume that parameters are age invariant (they depend only on the relative separation between the timing of the input and the measurement of skill) and that years of schooling (SY_{i2}) are contained within the vector of school inputs (SI_{i2}).⁶² This allows one to express the production function of skill as follows:

$$A_{i2} = HI'_{i2}\gamma_1 + HI'_{i1}\gamma_2 + SI'_{i2}\phi_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f'_i\lambda_{(2)} + \mu_{i0}\beta_{(2)} \quad (3.24)$$

As explained in the previous chapter, this can be regarded as a “cumulative” model where skill exhibited at the end of period 2 is expressed as a function of all relevant direct influences that took place until that moment. It is also possible to express A_{i2} as a function of lagged skill and period 2 influences only. For this, consider that period 1 skill can be written as:

$$A_{i1} = HI'_{i1}\gamma_1 + h_{i1}\varphi_1 + f'_i\lambda + \mu_{i0}\beta \quad (3.25)$$

The assumption of age-invariant parameters implies that γ_1 and φ_1 are the same in (3.24) and (3.25). In (3.25) they indicate the effect of period 1 educational and health inputs on period 1 skill. Also notice that parameters λ and β in (3.25) indicate the effect of predetermined direct influences and innate ability in period 1, respectively, while

⁶² The analysis can be extended to the more general case of age-dependent parameters at the cost of complicating notation with no effect on the main results.

parameters $\lambda_{(2)}$ and $\beta_{(2)}$ in (3.24) express the cumulative effect (until period 2) of this same pair of influences.

If we assume that the effect of inputs decays at a rate ρ (e.g. $\gamma_2 = \rho\gamma_1$) and subtract ρA_{i1} from (3.24) we obtain:⁶³

$$A_{i2} = \rho A_{i1} + HI'_{i2}\gamma_1 + SI'_{i2}\phi_1 + h_{i2}\varphi_1 + f'_i\lambda + \mu_{i0}\beta \quad (3.26)$$

where $\lambda = \lambda_{(2)} - \rho\lambda$ and $\beta = \beta_{(2)} - \rho\beta$. These parameters capture the contemporaneous (period 2) effect of predetermined direct influences and innate ability. Consistent with the notation introduced in the previous chapter, the expression given in (3.26) is known as a “value added” model.

Consistent estimation of the parameters involved in (3.24) is problematic because we seldom observe innate ability and all relevant inputs. A value added model can allow one to partially circumvent this problem if lagged skill is a sufficient statistic for assignment mechanisms that correlate with unobservable influences (e.g. if children end up sorted into different schools according to their pre-school skill). In this regard, several recent studies reviewed in Singh (2015) have shown that value added models can provide reliable estimates of the individual effects of skill inputs. A value added model, however, will not allow one to control for omitted period 2 inputs.

Hybrid models stand out as a popular empirical strategy to try to circumvent the problem of omitted inputs. As already explained, the objective is to control for omitted inputs using the arguments of their corresponding demand function. Quick inspection of Table 3.1 reveals that all the studies that have attempted to decompose the “explained” part of the gap have

⁶³ Appendix 3.1 presents more detail regarding how this assumption ensures that the model is no longer a function of period 1 inputs. It also presents the more general case of parameters that depend on child’s age. This should help clarify why the absence of period 1 inputs in (3.26) does not depend on the assumption of age-invariant parameters but only on the assumption that the effect of inputs decay at a rate equal to ρ .

relied on some form of hybrid specification. In fact, they all control for family or household characteristics that do not have a direct effect on skill but can influence it through the acquisition of educational home or school inputs.

The risk of obtaining biased estimates of the individual effects of inputs under different empirical specifications (including the value added and hybrid models presented here) has already been addressed in the literature (Todd and Wolpin, 2003, 2007). In the example that follows I will focus on another type of bias affecting linear gap decompositions that has not been acknowledged yet in the literature. As already mentioned, this has to do with the rules employed to assign variables into different categories.

For this, let us shift to the empirical versions of (3.24) and (3.25) assuming that cognitive skill is measured with random error through the scores obtained in some test: $T_{i2} = A_{i2} + \varepsilon_{i2}$; $E(\varepsilon_{i2}) = 0$, $Cov(\varepsilon_{i2}, A_{i2}) = 0$. Also assume there is a single unobserved input from each period, one belonging to the early childhood environment (HI_{i1}^U) and the other to the school environment (SI_{i2}^U).⁶⁴

This yields the following empirical version of the production function of skill:

$$T_{i2} = HI_{i2}'\gamma_1 + HI_{i1}'\gamma_2 + SI_{i2}'\phi_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f_i'\lambda_{(2)} + [HI_{i1}^U\gamma_2^U + SI_{i2}^U\phi_1^U + \mu_{i0}\beta_{(2)} + \varepsilon_{i2}] \quad (3.27)$$

This has a value-added representation given by:

$$T_{i2} = \rho T_{i1} + HI_{i2}'\gamma_1 + SI_{i2}'\phi_1 + h_{i2}\varphi_1 + f_i'\lambda + [SI_{i2}^U\phi_1^U + \mu_{i0}\beta + \varepsilon_{i2} - \rho\varepsilon_{i1}] \quad (3.28)$$

⁶⁴ The analysis can be extended to the more general case where we have several omitted inputs from both periods without affecting its main results.

The elements in brackets at the right hand side of (3.27) and (3.28) are contained in the corresponding error terms of both specifications.

Following the results of the model of family choice presented in the previous section, the demand functions of the inputs of skill (including those omitted) depend on predetermined household, family and child characteristics that influence skill directly (f_i) and other exogenous input determinants capturing differences in resources, prices, environments and preferences. Assume the latter are contained in a vector (z_i) and that demand functions can be expressed linearly. Accordingly, the demand functions for omitted inputs can be written as follows:

$$\begin{aligned} HI_{i1}^U &= z_i' \delta_1 + f_i' \kappa_1 + \tau_1 G_i + v_{i1} \\ SI_{i2}^U &= z_i' \delta_2 + f_i' \kappa_2 + \tau_2 G_i + v_{i2} \end{aligned} \quad (3.29)$$

Where v_{i1} and v_{i2} capture random shocks to the demand functions. Variable G_i denotes membership to the groups considered to define the gap in cognitive skill ($G_i = 1$ if the child belongs to group A and $G_i = 0$ if she belongs to group B). This indicator will typically have a role within z_i (and even f_i) as achievement gaps are usually defined in terms of children's ethnicity or their geographical domain (see Table 3.1). In other words, I am considering the fairly general case where the group indicator defining the achievement gap can be included among the arguments of the demand function of inputs.⁶⁵ In (3.29) I have considered this variable separate from z_i and f_i to ease the exposition of the decomposition strategies.

⁶⁵ For example, the geographical domain can be a relevant argument in a demand function insofar it controls for differences in exogenous environmental variables such as the general health status or the availability of educational services in the area (HE_{i1}, HE_{i2}, S_i in the model presented above). This is the theoretical motivation for the inclusion of the urban/rural group indicator in the hybrid specifications employed in Chapter 2. According to the exogenous nature of input determinants, group membership should not be part of the choices made by families during the period under analysis. Consistent with this, less than 5% of children considered in the sample employed in Chapter 2 (and the simulations that follow) changed their geographical domain between rounds.

If we replace (3.29) in (3.27) and collect terms, it is possible to build the following linear hybrid specification.

$$T_{i2} = HI'_{i2}\gamma_1 + HI'_{i1}\gamma_2 + SI'_{i2}\phi_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f'_i(\lambda_{(2)} + \gamma_2^U\kappa_1 + \phi_1^U\kappa_2) + z'_i(\gamma_2^U\delta_1 + \phi_1^U\delta_2) + (\tau_1\gamma_2^U + \tau_2\phi_1^U)G_i + [v_{i2}\phi_1^U + v_{i1}\gamma_2^U + \mu_{i0}\beta_{(2)} + \varepsilon_{i2}] \quad (3.30)$$

A similar exercise for the value added specification yields:

$$T_{i2} = \rho T_{i1} + HI'_{i2}\gamma_1 + SI'_{i2}\phi_1 + h_{i2}\varphi_1 + f'_i(\lambda + \phi_1^U\kappa_2) + z'_i\phi_1^U\delta_2 + \tau_2\phi_1^UG_i + [v_{i2}\phi_1^U + \mu_{i0}\beta + \varepsilon_{i2} - \rho\varepsilon_{i1}] \quad (3.31)$$

Expressions (3.30) and (3.31) lead to the following empirical specifications:

$$T_{i2} = HI'_{i2}\gamma_1 + HI'_{i1}\gamma_2 + SI'_{i2}\phi_1 + h_{i2}\varphi_1 + h_{i1}\varphi_2 + f'_i\pi + z'_i\psi + \theta G_i + e^H_{i2} \quad (3.32)$$

$$T_{i2} = \rho T_{i1} + HI'_{i2}\gamma_1 + SI'_{i2}\phi_1 + h_{i2}\varphi_1 + f'_i\tilde{\pi} + z'_i\tilde{\psi} + \tilde{\theta}G_i + e^{VA}_{i2} \quad (3.33)$$

The risk of obtaining biased estimates of the production function parameters from an OLS estimation of (3.32) or (3.33) arises from the potential correlation between observed inputs and the elements contained in the corresponding error terms of these equations. As already mentioned, this source of bias has been widely addressed in the literature and, in what follows, I will ignore it in order to focus on the bias that can arise from the rules employed to assign variables into categories as part of a gap decomposition exercise.

Let us define the achievement gap as the difference in expected skill between children belonging to groups A and B: $E(A_{i2}|A) - E(A_{i2}|B)$. Its empirical counterpart is given by: $\bar{T}_{A2} - \bar{T}_{B2}$, where upper bars indicate sample means. The inclusion of the group indicator in (3.32) and (3.33) ensures that an OLS regression passes through the mean of both groups. Thus, for the hybrid-cumulative specification we have:

$$\begin{aligned}\bar{T}_{A2} - \bar{T}_{B2} = & (\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1 + (\overline{HI}_{A1} - \overline{HI}_{B1})' \hat{\gamma}_2 + (\overline{SI}_{A2} - \overline{SI}_{B2})' \hat{\phi}_1 + (\bar{h}_{A2} - \bar{h}_{B2}) \hat{\phi}_1 \\ & + (\bar{h}_{A1} - \bar{h}_{B1}) \hat{\phi}_2 + (\bar{f}_A - \bar{f}_B)' \hat{\pi} + (\bar{z}_A - \bar{z}_B)' \hat{\psi} + \hat{\theta}\end{aligned}\quad (3.34)$$

And, for the hybrid-value added we have:

$$\begin{aligned}\bar{T}_{A2} - \bar{T}_{B2} = & (\bar{T}_{A1} - \bar{T}_{B1}) \hat{\rho} + (\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1 + (\overline{SI}_{A2} - \overline{SI}_{B2})' \hat{\phi}_1 + (\bar{h}_{A2} - \bar{h}_{B2}) \hat{\phi}_1 \\ & + (\bar{f}_A - \bar{f}_B)' \hat{\pi} + (\bar{z}_A - \bar{z}_B)' \hat{\psi} + \hat{\theta}\end{aligned}\quad (3.35)$$

If one seeks to decompose $\bar{T}_{A2} - \bar{T}_{B2}$, the feasible components will depend on the specification being used. Both specifications allow one to identify the contribution of observed inputs provided during period 2. The hybrid-cumulative specification also allows one to separately estimate the contribution of observed early childhood educational and health inputs. The hybrid-value added model, however, aggregates all early childhood influences into a single component $((\bar{T}_{A1} - \bar{T}_{B1}) \hat{\rho})$ containing the contribution of early childhood educational and health inputs, and also the contribution due to the early childhood effect of predetermined direct influences and innate ability.

In a similar fashion to most of the empirical work surveyed in Table 3.1, consider a researcher using a hybrid-cumulative specification and interested in comparing the contribution of “family influences” vs. the contribution of “school influences”. Another researcher using a hybrid-value added model seeks to compare the contribution of period 2 “family influences” vs. the contribution of “school influences”. For this, both need to devise a rule determining what constitutes a family influence. As documented in section 3.2, the rule commonly employed in the literature has been to assign all household, family and child characteristics (both inputs and input determinants) into this category. Consider that both researchers follow this standard rule and propose the decompositions presented in Table 3.3.

Table 3.3**Categories and variables included under the standard decomposition rule**

<i>Hybrid-cumulative</i>		<i>Hybrid-value added</i>	
Family influences ($\hat{H}$)	$(\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1$ $+ (\overline{HI}_{A1} - \overline{HI}_{B1})' \hat{\gamma}_2$ $+ (\bar{f}_A - \bar{f}_B)' \hat{\pi}$ $+ (\bar{z}_A - \bar{z}_B)' \hat{\psi}$	Period 2 family influences ($\hat{H}_2$)	$(\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1$ $+ (\bar{f}_A - \bar{f}_B)' \hat{\pi}$ $+ (\bar{z}_A - \bar{z}_B)' \hat{\psi}$
School inputs	$(\bar{SI}_{A2} - \bar{SI}_{B2})' \hat{\phi}_1$	School inputs	$(\bar{SI}_{A2} - \bar{SI}_{B2})' \hat{\phi}_1$
Health inputs	$(\bar{h}_{A2} - \bar{h}_{B2}) \hat{\phi}_1$ $+ (\bar{h}_{A1} - \bar{h}_{B1}) \hat{\phi}_2$	Period 2 health inputs	$(\bar{h}_{A2} - \bar{h}_{B2}) \hat{\phi}_1$
--	--	Past influences	$(\bar{T}_{A1} - \bar{T}_{B1}) \hat{\rho}$
Unexplained	$\hat{\theta}$	Unexplained	$\hat{\hat{\theta}}$

Consistent with the standard decomposition rule, the contributions of variables contained in f_i and z_i have been assigned to the category hosting family influences as they comprise characteristics that belong to this environment. Also notice that the contribution of the group indicator ($\hat{\theta}$ or $\hat{\hat{\theta}}$ depending on the specification) has been assigned to a component labelled “unexplained”. This ensures that these decompositions are consistent with the general OB approach employed in the literature. In fact, the two researchers of this example are using an OB decomposition where the “unexplained” part of the gap corresponds to the difference in coefficients measured with respect to a reference group built using a pooled regression which includes the group indicator (the decomposition labelled OB pooled* in Table 3.2). As explained in Section 3.2.1, the difference in coefficients between the two groups is given by the estimate of the parameters of the group indicator present in expressions (3.32) and (3.33).

To see the potential for bias in the two decompositions given above, let us focus on the estimated contribution of the two components related to “family influences”, starting with

$\hat{H}$. If we consider a sufficiently large sample and the parameter structure given in (3.30) it is not difficult to see that:

$$\begin{aligned}\hat{H} = & (\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1 + (\overline{HI}_{A1} - \overline{HI}_{B1})' \hat{\gamma}_2 + (\bar{f}_A - \bar{f}_B)' (\hat{\lambda}_{(2)} + \hat{\gamma}_2^U \hat{\kappa}_1 + \hat{\phi}_1^U \hat{\kappa}_2) \\ & + (\bar{z}_A - \bar{z}_B)' (\hat{\gamma}_2^U \hat{\delta}_1 + \hat{\phi}_1^U \hat{\delta}_2)\end{aligned}\tag{3.36}$$

From the demand functions of omitted inputs (see (3.29)) is possible to write:

$$\begin{aligned}\overline{HI}_{A1}^U - \overline{HI}_{B1}^U &= (\bar{z}_A - \bar{z}_B)' \hat{\delta}_1 + (\bar{f}_A - \bar{f}_B)' \hat{\kappa}_1 + \hat{\tau}_1 \\ \overline{SI}_{A2}^U - \overline{SI}_{B2}^U &= (\bar{z}_A - \bar{z}_B)' \hat{\delta}_2 + (\bar{f}_A - \bar{f}_B)' \hat{\kappa}_2 + \hat{\tau}_2\end{aligned}\tag{3.37}$$

Combining (3.36) and (3.37) we obtain:

$$\begin{aligned}\hat{H} = & (\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1 + (\overline{HI}_{A1} - \overline{HI}_{B1})' \hat{\gamma}_2 + (\bar{f}_A - \bar{f}_B)' \hat{\lambda}_{(2)} + \hat{\gamma}_2^U (\overline{HI}_{A1}^U - \overline{HI}_{B1}^U - \hat{\tau}_1) \\ & + \hat{\phi}_1^U (\overline{SI}_{A2}^U - \overline{SI}_{B2}^U - \hat{\tau}_2)\end{aligned}\tag{3.38}$$

The presence of bias is clear as this expression involves elements that belong to the contribution of school inputs $(\hat{\phi}_1^U (\overline{SI}_{A2}^U - \overline{SI}_{B2}^U - \hat{\tau}_2))$. Notice that the direct contribution of predetermined direct influences $((\bar{f}_A - \bar{f}_B)' \hat{\lambda}_{(2)})$ belongs to the “family influences” category. The contribution that operates through the demand of omitted inputs, however, does not fully belong to the “family influences” category because of the presence of omitted school inputs.

Consider a situation where the group indicator plays only a marginal role in the demand function of omitted inputs ($\tau_1 = \tau_2 = 0$). In this case, the estimated contribution of family influences $(\hat{H})$ would be able to account for the omitted early childhood input but at the

cost of overstating the importance of this category as it would also be including the contribution of the omitted school input.⁶⁶

Let us now analyse what lies behind the estimated contribution of period 2 family influences built using the hybrid-value added specification. A sufficiently large sample and the parameter structure given in (3.31) allow one to write:

$$\hat{H}_2 = (\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1 + (\bar{f}_A - \bar{f}_B)' (\hat{\lambda}_2 + \hat{\phi}_1^U \hat{k}_2) + (\bar{z}_A - \bar{z}_B)' \hat{\phi}_1^U \hat{\delta}_2 \quad (3.39)$$

Combining (3.37) and (3.39) we have:

$$\hat{H}_2 = (\overline{HI}_{A2} - \overline{HI}_{B2})' \hat{\gamma}_1 + (\bar{f}_A - \bar{f}_B)' \hat{\lambda}_2 + \hat{\phi}_1^U (\bar{SI}_{A2}^U - \bar{SI}_{B2}^U - \hat{\tau}_2) \quad (3.40)$$

In this case, the presence of a positive bias is even clearer. According to (3.40), the second researcher is able to fully account for the contribution of period 2 family influences (which include the period 2 effect of predetermined direct influences: $(\bar{f}_A - \bar{f}_B)' \hat{\lambda}_2$), but adds to this a positive element that corresponds to the contribution of school inputs $(\hat{\phi}_1^U (\bar{SI}_{A2}^U - \bar{SI}_{B2}^U - \hat{\tau}_2))$.

This analysis illustrates how, unless we are able to claim the all of the omitted inputs belong only to a certain category, including variables that reflect predetermined direct influences or input determinants into this category will likely lead to a positive bias in its estimated contribution. In particular, it has exposed how decompositions based on the standard rule of assigning all available household, family and child characteristics into a single category can lead to an overstatement of the importance of influences related to the home or family environment *vis-à-vis* that of school inputs. This is especially significant when data on

⁶⁶ Notice I am assuming that all inputs have a positive effect on skill (i.e. that the parameters in the production function of skill are all positive) and that group A exhibits an advantage with respect to group B.

school characteristics are not particularly rich, so omitted inputs include several influences that belong to the school environment.

Based on this, in what follows I propose an alternative decomposition strategy. This strategy acknowledges the difference and relations between skill inputs and skill input determinants, and corresponds to the one employed for the hybrid models used in Chapter 2. Accordingly, it is based on assigning the contribution of all predetermined direct influences and exogenous input determinants into an especial variable category hosting omitted inputs. In what follows, I briefly describe the strategy considering the notation used in this section and explain why it can be regarded as a special type of OB decomposition.

3.4.2 An alternative decomposition strategy: Oaxaca-Blinder with a twist

The strategy takes into account the parameter structure behind the empirical specifications of the hybrid and hybrid-value added models. This structure stems from the fact that predetermined direct influences (f_i), exogenous input determinants (z_i) and the group indicator (G_i) control for all omitted inputs through their demand equations. As described in equations (3.30) and (3.32), parameters accompanying f_i , z_i and G_i in the hybrid specification considered for the example given above are as follows: $\pi = \lambda_{(2)} + \gamma_2^U \kappa_1 + \phi_1^U \kappa_2$; $\psi = \gamma_2^U \delta_1 + \phi_1^U \delta_2$; and $\theta = \tau_1 \gamma_2^U + \tau_2 \phi_1^U$. Parameters accompanying these same variables in the hybrid-value added specification are as follows (see equations (3.31) and (3.33)): $\tilde{\pi} = \lambda + \phi_1^U \kappa_2$; $\tilde{\psi} = \phi_1^U \delta_2$; and $\tilde{\theta} = \tau_2 \phi_1^U$.

As discussed when presenting the urban/rural decompositions performed in the previous chapter, an important implication of this parameter structure is that it will not be possible to separately identify the direct and indirect effects of predetermined direct influences (f_i) unless we impose further restrictions. In addition, strong assumptions are also required to claim that omitted inputs belong only to either the family or school environment in order to

assign the contribution of all the arguments belonging to demand functions to one of these categories.

Based on the above, the decomposition strategy proposed here will assign the contribution of all variables contained in f_i and z_i and the indicator function (G_i) into a special category hosting the contribution of predetermined direct influences and omitted inputs in general. In the particular case of the hybrid-value added plus specification, this joint contribution will consider only period 2 predetermined direct influences and period 2 omitted inputs. Table 3.4 summarizes the categories proposed based on the contributions given in (3.34) and (3.35).

Table 3.4
Categories and variables included under the alternative decomposition strategy

<i>Hybrid-cumulative</i>		<i>Hybrid-value added</i>	
Early childhood and home inputs	$(\bar{H}I_{A2} - \bar{H}I_{B2})'\hat{\gamma}_1 + (\bar{H}I_{A1} - \bar{H}I_{B1})'\hat{\gamma}_2$	Period 2 home inputs	$(\bar{H}I_{A2} - \bar{H}I_{B2})'\hat{\gamma}_1$
Health inputs	$(\bar{h}_{A2} - \bar{h}_{B2})\hat{\phi}_1 + (\bar{h}_{A1} - \bar{h}_{B1})\hat{\phi}_2$	Period 2 health inputs	$(\bar{h}_{A2} - \bar{h}_{B2})\hat{\phi}_1$
School inputs	$(\bar{S}I_{A2} - \bar{S}I_{B2})'\hat{\phi}_1$	School inputs	$(\bar{S}I_{A2} - \bar{S}I_{B2})'\hat{\phi}_1$
Predetermined direct influences and omitted inputs	$(\bar{f}_A - \bar{f}_B)'\hat{\pi} + (\bar{z}_A - \bar{z}_B)'\hat{\psi} + \hat{\theta}$	Period 2 predetermined direct influences and period 2 omitted inputs	$(\bar{f}_A - \bar{f}_B)'\hat{\pi} + (\bar{z}_A - \bar{z}_B)'\hat{\psi} + \hat{\theta}$
--	--	Past influences	$(\bar{T}_{A1} - \bar{T}_{B1})\hat{\rho}$

At this point is worth recalling from the discussion in Chapter 2, that the parameter structure described above allows for a simple test for omitted inputs. In particular, rejection of the null hypothesis $(\bar{z}_A - \bar{z}_B)'\psi + \theta = 0$ in the hybrid-cumulative specification implies the presence of at least one omitted input. Rejection of the null $(\bar{z}_A - \bar{z}_B)'\tilde{\psi} + \tilde{\theta} = 0$ in the hybrid-value added specification implies the presence of at least one omitted period 2 input.

From the discussion in section 3.2 is clear that there are different ways to implement the OB decomposition and that there are several possible interpretations for the “unexplained” part of the gap. One can choose between a “threefold” or a “twofold” decomposition, and one needs to decide which will be the reference group used to measure the difference in coefficients that produces the “unexplained” part of the gap. This “unexplained” part, in turn, can be interpreted as actually capturing a difference in returns to inputs or the presence of omitted inputs.

The decomposition strategy described here can be considered as a special case of “twofold” OB decomposition. In fact, inclusion of the group indicator G_i in the hybrid models described above ensures that I am using the same coefficient estimates as those required to build the OB decomposition that uses the results of a pooled regression as a reference (OB pooled*; see Table 3.2). In addition, this can be regarded as an OB decomposition where the “unexplained” part of the gap is interpreted as capturing the contribution of omitted inputs. Notice this means that I am not imposing the restriction of equal coefficients in both groups when estimating the empirical specifications. What I am restricting is the interpretation of this coefficient difference (or “unexplained” part) to mean the presence of omitted inputs.

The distinctive feature of this decomposition strategy with respect to those surveyed in Table 3.1 is that it is based on the results of a model that postulates a relation between cognitive skill and its inputs, and describes how families’ choices determine these inputs. In terms of an OB approach, this prevents arbitrary choices of reference group and interpretation of the “unexplained” part of the gap. It also makes explicit the difference between inputs and input determinants which, in turn, prevents the use of rules that can introduce bias by assigning the contribution of one category to another. In fact, this is the reason why, as opposed to a more standard OB decomposition, in this strategy the group indicator is not the only variable accounting for omitted inputs. Predetermined direct

influences and exogenous input determinants are also included in the category hosting omitted inputs as they all have a role as arguments in their demand functions.

This feature of the decomposition strategy is related to the message conveyed by Neumark's analysis (Neumark, 1988). He considered the use of different reference groups when estimating the contribution of "discrimination" to a wage gap through an OB decomposition. He traced back the choice of a certain reference group to a particular assumption regarding firms' discriminatory behaviour.⁶⁷ In similar fashion, if one seeks to measure the contribution of different types of inputs to a certain cognitive skill gap by means of a particular decomposition strategy, one needs to be aware of the assumptions in terms of family behaviour that allow one to recover these contributions.⁶⁸ In sum, to avoid arbitrary choices in terms of components and interpretations, the choice of decomposition strategy should not be made without consideration of the underlying behavioural assumptions that allow one to identify the contributions of interest.

3.5 An illustration using Peruvian data

In what follows I will empirically illustrate the main issues discussed so far. Namely: (i) that assigning all available household, family and child characteristics into a single category will likely lead to an overstatement of the importance of direct influences originated at home *vis-à-vis* that of school inputs, especially when one lacks rich information on school characteristics; (ii) that the use of school fixed effects to account for the contribution of school inputs can lead to an overstatement of the importance of these influences, especially when one lacks information on early childhood inputs or lagged skill and schools are highly

⁶⁷ In Neumark (1988), the author showed that choosing the advantaged group as a reference implied assuming that the firm only discriminates against the disadvantaged group ("pure discrimination"). Choosing the disadvantaged group as a reference implies that the firm only prefers the advantaged group ("pure nepotism").

⁶⁸ For example, assuming that family choices play no role in determining the characteristics of schools hosting their children would allow one to assign all demand function arguments to the early childhood and home inputs category. Doing this would resemble the standard decomposition rule.

segregated; and (iii) that the alternative decomposition strategy proposed here is less prone to biases than those employed so far in the literature.

For this, I will decompose the gap in cognitive skill observed between urban and rural 8-year-old children in Peru. I will first build a “full information” decomposition relying on an unusually rich dataset that contains abundant information on school inputs and longitudinal information on cognitive achievement. This will be based on the components of the alternative decomposition strategy as described in Table 3.4. The objective is to verify that the school inputs considered do make a significant contribution to the gap under analysis.

I will then exclude school input information and perform three additional decompositions. The first will be based on the components of the alternative decomposition strategy and will serve to determine whether predetermined direct influences and exogenous input determinants pick-up the contribution of the omitted school inputs as predicted by the model described in section 3.3. The second will be based on the components under the standard decomposition rule and will be used to verify that this decomposition introduces a positive bias in the estimated contribution of family and household influences. The third will include school fixed effects to account for school inputs and will be compared against the “full information” decomposition to assess if the fixed effects estimation tends to overstate the importance of school inputs.

3.5.1 Data sources and variables

The data for this analysis is the same as the one employed for the urban/rural decompositions performed in Chapter 2. Here I only present a brief overview and the reader is referred to Section 2.4.3 for a more detailed description.

The following exercise will consider the information contained in the first three rounds of the child and household surveys, as well as the school survey, focusing on the Younger Cohort of the Young Lives Study in Peru. The simulations will be based on two different samples. The first considers all children that have cognitive test scores for rounds 2 and 3, and attend a school included in the school survey (487 children in 124 schools). The second sample considers all children that have cognitive test scores for rounds 2 and 3 (1,561 children).

Following the analytical framework described in section 3.3, period 1 variables will correspond to influences relevant from birth and up to age 5, and period 2 variables will correspond to influences relevant between ages 5 and 8. Accordingly, period 1 variables will be provided by rounds 1 and 2, while period 2 variables will be provided by round 3. Influences captured in the school survey (collected two years after round 3) will be assumed to be the same as those present in period 2.⁶⁹

The measures of cognitive achievement employed in this analysis are the standardized test scores obtained in the Peabody Picture Vocabulary Test (PPVT). Rounds 1, 2 and 3 of the household and child survey contain rich information on household and caregiver characteristics. Following the analytical framework, these were grouped into early childhood and home inputs, health inputs, predetermined direct influences, and input determinants. Following the classifications proposed in Guerrero et al. (2012), the information contained in the school survey was grouped into six categories: (i) school size, organization and timetable; (ii) infrastructure; (iii) climate⁷⁰; (iv) learning activities and

⁶⁹ I am assuming that school characteristics have not changed significantly during the two year period that separates round 3 from the school survey and that the child has remained in the same school since her enrolment in first grade (at age 6) until the moment in which the school survey was collected (at age 10). Accordingly, the administrative data collected from the schools included in the survey reveals that school switching is a rare phenomenon. On average, only 2% of students enrolled in primary education changed school each year between 2009 and 2010.

⁷⁰ Variables in this category include teachers' perception of the relations among students and between students and teachers, and of the problems and difficulties encountered during the school year.

materials; (v) teacher characteristics; and (vi) school responsiveness⁷¹. Tables 2.5 and 2.6 (in Section 2.4.3) provide the complete list of variables as well as their descriptive statistics and urban/rural differences.

3.5.2 Decomposition results and discussion

In this section I present and discuss the results obtained after estimating the contributions indicated in Tables 3.3 and 3.4.⁷² In particular, Table 3.7 presents the estimated contributions of each component considering: (i) the alternative decomposition strategy and all the information provided by the school survey (the “full information” decomposition; see panel A); (ii) the alternative decomposition strategy excluding the inputs provided by the school survey⁷³ and using the same sample as in (i) (see panel B); (iii) the alternative decomposition strategy excluding the inputs provided by the school survey and using the complete sample of children that register a PPVT score in rounds 2 and 3 (see panel C); (iv) the standard decomposition rule excluding the inputs provided by the school survey and using the complete sample of children (see panel D); and (v) the alternative decomposition strategy excluding the inputs provided by the school survey but using school fixed effects to account for the contribution of school inputs (see Panel E).

Figure 3.1 shows the same point estimates accompanied by 95% confidence intervals. It also presents the statistic and corresponding p-value of the test of omitted inputs described in the previous section. Recall that this statistic provides an estimate of the contribution of exogenous input determinants to the gap under analysis. Appendix 3.2 presents coefficient estimates for the variables involved in all the specifications. I also tested the presence of

⁷¹ Variables within this category indicate whether or not the school provides support for students lagging behind or at risk of dropping out.

⁷² The hybrid-value added models were estimated including also the period 1 inputs available. This does not alter the interpretation of its coefficients and is a less restrictive specification as it relaxes the assumption requiring that the effects of period 1 inputs decay at a rate ρ . Consistent with the logic of a value added specification, the contributions of included period 1 inputs were assigned to the “past influences” category.

⁷³ Ignoring the information contained in the school survey implies that the only school inputs considered are years of schooling and time spent at school.

dynamic complementarity by introducing interactions between school quality variables and past cognitive achievement in the value added specification. None of the interaction terms were found to be statistically significant, both independently and jointly (see Table 3C in Appendix 3.2).

The first set of results reveals that school inputs have a significant contribution of around 35% to the cognitive skill gap observed at age 8 between urban and rural children. To account for this contribution I am reporting the estimate provided by the hybrid-value added specification. As mentioned above, several recent empirical studies have shown that this specification can provide reliable estimates of the effect of contemporaneous influences on skill, revealing that lagged test scores are a sufficient statistic for input assignment mechanisms that correlate with unobservables such as omitted past inputs and innate ability. Consistent with this, the hybrid-cumulative model, which can control for omitted inputs but retains the full cumulate effect of unobserved innate ability, shows a somewhat larger estimated contribution for school inputs.

Results reported in panel A of Figure 3.1 show that exogenous input determinants in the hybrid-cumulative specification have a significant contribution (22%; $p < 0.05$), indicating the presence of omitted inputs being controlled for through their demand equations. It is worth noticing that in the hybrid-value added model it is no longer possible to reject the null that input determinants are non-significant. Since in this model we are controlling for all period 1 inputs (both observed and omitted), this evidence is consistent with the absence of period 2 omitted inputs.⁷⁴

⁷⁴ Failure to reject the null $(\bar{z}_A - \bar{z}_B)' \tilde{\psi} + \tilde{\theta} = 0$ does not directly imply the absence of omitted inputs. If predetermined direct influences are a sufficient statistic in the demand equation of omitted inputs, the null $(\bar{z}_A - \bar{z}_B)' \tilde{\psi} + \tilde{\theta} = 0$ will not be rejected even under the presence of omitted inputs. In this case, however, we cannot say that predetermined direct influences are a sufficient statistic in the demand equation of omitted inputs because exogenous input determinants do have a significant contribution in the hybrid-cumulative model.

Table 3.7
Normalized contributions to the urban/rural gap in cognitive skill at age 8
(% of urban-rural gap)

Hybrid-cumulative		Hybrid-value added	
(A) “Full information” decomposition: includes all inputs from the school survey			
Early childhood and home inputs	0.129*** (0.049)	Period 2 home inputs	0.030 (0.041)
School inputs	0.479*** (0.086)	School inputs	0.348*** (0.081)
Health inputs	0.063** (0.029)	Period 2 health inputs	0.041 (0.029)
--	--	Past influences	0.368*** (0.081)
Predetermined direct influences and omitted inputs	0.328*** (0.089)	Period 2 predetermined direct influences and period 2 omitted inputs	0.214** (0.082)
(B) Alternative decomposition strategy excluding inputs from the school survey (reduced sample)			
Early childhood and home inputs	0.216*** (0.061)	Period 2 home inputs	0.071* (0.041)
School inputs	0.062*** (0.016)	School inputs	0.045*** (0.014)
Health inputs	0.077*** (0.029)	Period 2 health inputs	0.044 (0.031)
--	--	Past influences	0.416*** (0.098)
Predetermined direct influences and omitted inputs	0.645*** (0.075)	Period 2 predetermined direct influences and period 2 omitted inputs	0.424*** (0.095)
(C) Alternative decomposition strategy excluding inputs from the school survey (complete sample)			
Early childhood and home inputs	0.200*** (0.024)	Period 2 home inputs	0.082*** (0.022)
School inputs	0.051*** (0.009)	School inputs	0.027*** (0.007)
Health inputs	0.060*** (0.023)	Period 2 health inputs	0.025** (0.012)
--	--	Past influences	0.419*** (0.032)
Predetermined direct influences and omitted inputs	0.689*** (0.057)	Period 2 predetermined direct influences and period 2 omitted inputs	0.447*** (0.055)

<i>Hybrid-cumulative</i>		<i>Hybrid-value added</i>	
<i>(D) Standard decomposition rule excluding inputs from the school survey (complete sample)</i>			
Family influences	0.470*** (0.046)	Period 2 family influences	0.302*** (0.045)
School inputs	0.051*** (0.009)	School inputs	0.027*** (0.007)
Health inputs	0.060** (0.023)	Period 2 health inputs	0.025** (0.012)
--	--	Past influences	0.419*** (0.032)
Unexplained	0.419*** (0.079)	Unexplained	0.228*** (0.077)
<i>(E) School fixed effects as school inputs (complete sample)</i>			
Early childhood and home inputs	0.117 (0.110)	Period 2 home inputs	0.030 (0.072)
School inputs	0.720*** (0.106)	School inputs	0.524*** (0.103)
Health inputs	0.068** (0.027)	Period 2 health inputs	0.042 (0.033)
--	--	Past influences	0.388*** (0.078)
Predetermined direct influences and omitted inputs	0.094 (0.116)	Period 2 predetermined direct influences and period 2 omitted inputs	0.016 (0.111)

Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Figure 3.1
Normalized
contributions to
the urban/rural
gap in cognitive
skill at age 8
(point estimates
and 95%
confidence
intervals)




Figure 3.1
Normalized
contributions to
the urban/rural
gap in cognitive
skill at age 8
(point estimates
and 95%
confidence
intervals)





Figure 3.1
Normalized
contributions to
the urban/rural
gap in cognitive
skill at age 8
(point estimates
and 95%
confidence
intervals)



Notes:
“EC&HomeInputs” refers to early childhood and home inputs.
“Pred&Omitted” refers to predetermined direct influences and omitted inputs.
“P2HomeInputs” refers to period 2 home inputs.
“P2HealthInputs” refers to period 2 health inputs.
“PastInf” refers to past influences.
“P2Pred&Omitted” refers to period 2 predetermined direct influences and period 2 omitted inputs.
“P2FamilyInf” refers to period 2 family influences.

If the contribution of omitted inputs is captured by their demand equations when exogenous input determinants are included in the regression, omission of significant school inputs should increase the contribution of the “predetermined direct influences and omitted inputs” category. Results presented in Panel B of Table 3.7 and of Figure 3.1 are consistent with this. In particular, the contribution of the category hosting omitted inputs in the hybrid-cumulative specification grows twice as large when school survey information is omitted (from 33% in the “full information” decomposition up to 65%). There is also strong evidence of the presence of omitted inputs as the contribution of exogenous input determinants in the hybrid cumulative model is now 51% (it was 22% in the “full information” decomposition) and highly significant ($p < 0.00$).

There is also a significant increase in the contribution of the category hosting omitted inputs in the value added specification (from 21% in the “full information” decomposition up to 42% after ignoring school inputs provided by the school survey). Importantly, the contribution of exogenous input determinants now remains significant (30%; $p < 0.00$) which implies we cannot accept the null $(\bar{z}_A - \bar{z}_B)' \tilde{\psi} + \tilde{\theta} = 0$ (see Panel B in Figure 3.1). This result, which differs from the one obtained with the complete set of data, confirms there are still relevant period 2 influences omitted. This is consistent with the fact that we are intentionally omitting school inputs.

It is also worth mentioning that, consistent with the fact that exogenous input determinants and predetermined direct influences are controlling for the omitted school inputs, we observe only a small increase in the estimated contribution of the categories hosting home inputs. This increase, which could be regarded as a bias, remains well within standard errors in both specifications (compare panels A and B in Figure 3.1).

Results presented in Panel C of Table 3.7 and of Figure 3.1 reveal that the results just discussed are robust to considering the entire sample of children with complete PPVT scores and not just those attending schools included in the school survey. This should mitigate concerns regarding potential selection bias in the sample used for the preceding analysis. The use of a larger sample also adds precision to the results discussed in the previous paragraphs.

The fact that the variables included in the “predetermined direct influences and omitted inputs” category are capturing the contribution of omitted school inputs implies that assigning these variables to a category that does not correspond to the school environment will generate a bias. This is precisely what happens if we omit school inputs provided by the school survey and use the standard decomposition rule. All family, household and child characteristics are assigned to a single category (“family influences”) while the remaining urban-rural indicator is used to capture the “unexplained” part of the cognitive skill gap. The new category hosting “family influences” has a contribution between 20 and 25 percentage points larger than the one capturing early childhood and home inputs in the alternative decomposition strategy (compare panels C and D in Figure 3.1). We know at least part of this additional contribution is a bias because at least part of it belongs to the school environment through the school inputs we are intentionally omitting.

It is worth noticing that there is also an important difference in the way one would interpret the portion of the gap that cannot be explained by the observed influences. To see this, let us refer to the decompositions based on the value added model in panels C and D of Figure 3.1. Under the standard decomposition rule (panel D) and in absence of structure regarding the source of the “unexplained” part of the gap, one would conclude that past achievement and the available family and school influences explain nearly 80% of the cognitive skill gap and

that most of this contribution has to do with influences that belong to the home environment. In addition, one could say that the remaining 20% of the gap remains unexplained and that this could be due to the omission of relevant skill determinants or to the fact that urban and rural children transform inputs into test scores differently (i.e. they exhibit different returns for a given set of inputs).⁷⁵

If we follow the alternative decomposition strategy (panel C), however, one would notice that a contribution similar to that of past influences (around 40%) can be attributed to period 2 predetermined direct influences and omitted inputs. This estimate is not only closer to the contribution of the school inputs we are intentionally omitting (which is around 35%) but its interpretation is much more informative as it explicitly acknowledges the presence of omitted inputs that belong to the second period.

Finally, panel E in Table 3.7 and Figure 3.1 present decomposition results using school fixed effects instead of the inputs contained in the school survey. A comparison against the “full information” decomposition (see Panel A) reveals a tendency to overstate the contribution of school inputs, especially in the absence of information on past achievement (i.e. in the hybrid-cumulative model). In fact, in the hybrid-cumulative model is quite clear that the school fixed effect has absorbed missing inputs from the early childhood period leaving the “predetermined direct influences and omitted inputs” category with an insignificant contribution, as opposed to what happened in the “full information” decomposition.

⁷⁵ Notice that for the standard rule I have employed an OB pooled* decomposition. If we use an OB pooled decomposition (where reference coefficients are provided by a pooled regression that excludes the group indicator), we would obtain an even smaller estimate for the “unexplained” part of the gap (13%), which implies a larger risk of overstating the contribution of home influences.

3.6 Concluding remarks

Linear decompositions based on the Oaxaca-Blinder technique are a fairly common way of attempting an estimate of the contribution of two or more categories of variables to the difference in cognitive achievement between children of different socioeconomic backgrounds. Two prominent examples of these categories are family and school influences.

In this chapter, I have argued that performing such decompositions in the absence of a framework postulating how cognitive skill is accumulated and how its inputs are determined, can be problematic in several ways. In particular, absence of this framework can lead one to overlook the difference between skill inputs and skill input determinants and to make arbitrary choices in terms of decomposition strategy and interpretation of its components. This, in turn, can lead to biases in the estimated contributions and to misleading policy implications.

This analysis has reviewed several studies using data from developing countries⁷⁶ and has found no consensus regarding the specific Oaxaca-Blinder strategy to use and how to interpret the “unexplained” part of the gap. These studies also exhibit a tendency to group all observed family, household and child characteristics in a category different from the school environment (see, for example, McEwan and Marshall (2004) or Hernandez-Zavala et al. (2006)). I have argued this can lead to an overstatement of the importance of family and household influences because several of these characteristics can reflect omitted inputs

⁷⁶ The analysis presented in Krieg and Storer (2006) is an interesting example from the developed world. The authors aim at determining how much of the difference between high and low-performing schools in the state of Washington can be attributed to characteristics that are beyond the control of school administrators. They used an OB twofold decomposition and found that a significant proportion of the gap can be attributed to characteristics of students and their families. Accordingly, they concluded arguing against penalizing poor performing schools. Interestingly, these authors acknowledge that student characteristics can pick-up the influence of omitted school inputs if sorting into different schools correlates with students’ socioeconomic status. The authors called this the “Tiebout effect” based on Tiebout’s model where agents choose public goods by deciding where to live (Tiebout, 1956). In a similar fashion, the model with endogenous school quality proposed here also gives a role to families’ choices as a determinant of school inputs.

that belong to the school environment. In addition, school fixed effects can pick-up the contribution of omitted inputs that belong to the home environment, especially in highly segregated school systems. There are, therefore, reasons to doubt the significant school contributions found in those studies that have used this empirical strategy (see, for example, McEwan (2004) or McEwan and Trowbridge (2007)).

Based on Glewwe and Miguel (2008), I developed a simple model explaining the skill formation process and how its inputs, including school characteristics, are determined by families' choices. I then used these insights to illustrate, analytically, the potential biases that can be introduced by the rule of grouping all family, household and child characteristics into a single category. I also used the results of the model to justify the categories proposed for an alternative decomposition strategy that aims at being less susceptible to biases that those used so far in the literature.

This alternative strategy uses the coefficient estimates of a particular type of Oaxaca-Blinder decomposition, but arranges the contribution of individual variables considering that predetermined family, household and child characteristics belong to the demand functions of inputs. These variables are, therefore, grouped in a special category hosting omitted inputs that resembles the "unexplained" part of the gap in a more conventional Oaxaca-Blinder decomposition.

Finally, I illustrated empirically the main issues discussed in the analysis by decomposing the gap in cognitive achievement between urban and rural 8-year-old children in Peru. I relied on an unusually rich dataset containing comprehensive information on school characteristics and longitudinal information on skill that allows one to control for early childhood inputs and innate ability. This provided a fairly reliable estimate of the contribution of school inputs. I then intentionally omitted information on school inputs and

found that: (i) their contribution is picked-up by the predetermined family, household and child variables included, as predicted by the framework that understands these as arguments of their demand functions; (ii) assigning these exogenous input determinants into a single “family influences” category will lead to an overstatement of their contribution; (iii) the alternative decomposition strategy proposed here is less prone to this kind of bias and provides a more accurate interpretation of the “unexplained” part of the gap; and (iv) the use of school fixed effects leads to an overstatement of the contribution of school inputs, especially when information on early childhood influences is lacking.

This analysis has provided a stronger and more formal foundation for the decomposition strategy employed in Chapter 2. In addition, if one seeks to perform a linear decomposition of an achievement gap, this analysis suggests that one needs to exercise caution when classifying the potential covariates. In particular, it is always advisable to determine which of the available variables better reflect an input of skill and which better reflect an input determinant (i.e. an argument in the demand function of the inputs of skill). Based on this, one can then assess which are the assumptions required for input determinants to control for omitted inputs that belong only to a certain category. For example, using predetermined family, household and child characteristics to control only for omitted home influences implies there are no omitted school inputs or that family decisions have not a significant role in determining the quality of the school environment. If these assumptions are not plausible considering the data in hand and the education system under analysis, one can rely on the decomposition strategy proposed here.

Table 3.1**Studies that have attempted a linear decomposition of the observed cognitive achievement gap between two groups of children**

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
1	McEwan and Trowbridge (2007) Guatemala	Test scores in language and maths Indigenous vs. non-indigenous children in grades 3 and 6	Regressions on contemporaneous predictors. OB twofold; reference coefficients from pooled regression including group indicator (pooled*) Components: endowments ("explained"), group indicator ("unexplained")	C1: Family variables C2: Quality of schools	C1: Parental education, presence of books, television viewing, child sex C2: School fixed effects PROENERE survey (2001)	> Explained component (language- mathematics) Grade 3: 71%-77% Grade 6: 55%-68% > Categories (language-mathematics) Grade 3: C1= 6%-8%; C2 = 65%-69% Grade 6: C1 = 5%-3%; C2 = 50%-66% Remarks: Between 50-69% of the gap is explained by the varying quality of schools that are attended by indigenous and non- indigenous children.
2	McEwan (2004) Bolivia Chile	Test scores in language and maths Indigenous vs. non-indigenous children in grades 3 and 6 (Bolivia) and 4 and 8 (Chile)	Regressions on contemporaneous predictors OB twofold; reference coefficients from pooled regression including group indicator (pooled*) Components: endowments ("explained"), group indicator ("unexplained")	C1: Family variables C2: School variables C3: Classroom variables	C1: Parental education, access to basic services (Bolivia), income (Chile), presence of books at home (Chile) C2: School fixed effects C3: Class fixed effects Bolivia: SIMECAL (1997) Chile: SIMCE (1997-1999)	> Explained component Between 80% and 90% > Categories C1 between 20% and 40% C2 and C3 between 50% and 70% Remarks: Between 50% and 70% of the gaps are attributable to differences in the quality of schools and classrooms. The gap may be the result of an unequal distribution of school and classroom resources, but also of an unequal distribution of peer-group characteristics.

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
3	Ramos et al. (2012) Colombia	Test scores in mathematics, science, reading Urban vs. rural students	Regressions on contemporaneous predictors OB twofold; reference coefficients from rural group Components: endowments ("explained"), coefficients ("unexplained")	C1: Individual characteristics C2: Family characteristics C3: School characteristics	C1: Child's age, child's sex C2: Parental education, second generation migrant, home language, presence of books at home C3: School is private/public, size, students per teacher, mean socioeconomic level of peer group ("peer effects") PISA survey (2006 and 2009)	> Explained component 100% of gap ("unexplained" component has negative contribution) > Categories C3 between 75% and 83% (mostly explained by the mean socioeconomic level of peer group) Remarks: Most of the rural-urban school differential is related to family characteristics. Policy implications: measures aimed at improving the general educational situation and conditions in the family.
4	McEwan and Marshall (2004) Cuba Mexico	Test scores in language and maths Mexican vs. Cuban children in grades 3 and 4	Regressions on contemporaneous predictors OB twofold; reference coefficients from low- achieving country (Mexico) and high- achieving country (Cuba) Components: endowments ("explained"), coefficients ("unexplained")	C1: Student and family variables C2: Peer variables C3: School and teacher variables	C1: Parental education, access to books, child's sex C2: Class average of parental education and access to books, climate in class C3: Grade, access to preschool, access to textbooks, teacher training and qualifications, access to materials, classroom conditions, school is urban/rural UNESCO Assessment of Student Achievement (1997)	> Explained component Between 11% and 31% > Categories C1 between 3% and 10% C2 between 14% and 23% C3 less than 3% Remarks: No more than 30% of the difference can be explained by differences in family, peer and school characteristics. Peer group and family variables contribute the largest portion to the explained component. Observed school and teacher variables are inconsistently linked to achievement.

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
5	Hernandez-Zavala et al. (2006) Guatemala Mexico Peru	Test scores in language and maths Indigenous vs. non-indigenous children between ages 8-10	Regressions on contemporaneous predictors OB twofold; reference coefficients from non-indigenous group Components: endowments ("explained"), coefficients ("unexplained")	C1: Family and child inputs C2: School inputs	C1: Parental education, access to basic services (Mexico, Guatemala), household assets (Guatemala), books at home (Peru), child's sex, child's grade, child repeated a grade, child works, child attended preschool C2: School is private/public; urban/rural (Peru, Guatemala), teacher experience, classroom condition (Peru and Mexico), access to textbooks (Guatemala), pupil-teacher ratio (Peru, Guatemala) Peru: First Comparative International Study on Language, Mathematics and Associated Factors (1997) Guatemala: "Laboratorio Guatemala" (2002) Mexico: National Standards (2001)	> Explained component (language-mathematics) Guatemala: 41%-55% Mexico: 75%-68% Peru: 70%-66% > Categories (language-mathematics) Guatemala: C1 = 23%-33%; C2 = 17%-23% Mexico: C1 = 67%-75%; C2 = 0% Peru: C1= 38%-41%; C2 25%-32% Remarks: Family variables contribute more than school variables to the overall explained component. Policy implications: effective bilingual education, compensatory education programs, choice of school and increased autonomy.

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
6	Arteaga and Glewwe (2014) Peru	Test scores in PPVT and maths Indigenous vs. non-indigenous children at ages 8 and 5	Regressions on past and contemporaneous predictors. Separate models for children at 8 and 5 years of age OB threefold; reference group is the indigenous group Components: endowments, coefficients ("heterogeneous effects"), interaction	C1: Household and child characteristics C2: Community characteristics	C1: Household expenditure, parental education, school expenditure, months in day care, months breastfeeding, prenatal visits, child did homework with parents, child played with parents, child's age, child's sex, child's nutritional status. C2: Community fixed effects Young Lives Study, rounds 2 (2006) and 3 (2009)	Most significant contributors: > By 5 years of age; PPVT Endowments of C2 = 67% > By 8 years of age; PPVT and maths Endowments of C1 > 80% (especially in parental education). * The contribution of interaction terms is not reported Remarks: By age 8 the importance of community characteristics recedes and household and child characteristics play the major role. Policy implication: Increase indigenous children's years of education as they will be household heads in the future.

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
7	Barrera- Osorio et al. (2011) Indonesia	Test scores in mathematics 2006 vs. 2003 results	Regressions on contemporaneous predictors OB twofold; (a) reference coefficients from 2006 group; (b) reference coefficients from pooled regression including group indicator Components: (a) endowments ("explained"), coefficients ("unexplained"); (b) endowments ("explained"), group indicator ("unexplained")	C1: Institutions / schools C2: Student characteristics C3: Family background	C1: School determines pedagogy, adequate supply of teachers, private-public, urban- rural, % repeating grade C2: grade, child's age, child's sex C3: Parental education, books present at home, access to computer, mother tongue PISA survey (2003 and 2006)	> Unexplained component Between 63% and 92% > Categories Most of the change in returns was in C2 Remarks: Most of the test score increase between 2003 and 2006 was due to changes in the returns to the characteristics rather than due to changes in the characteristics themselves.

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
8	Zhang and Lee (2011) Korea Japan OECD countries	Test scores in mathematics, science, reading Specific OECD countries vs. OECD average	Regressions on contemporaneous predictors Resembles OB twofold; gaps are measured against OECD average using predictor means and coefficients from a regression containing all OECD countries Components: endowments ("explained"), coefficients ("unexplained")	C1: Individual characteristics C2: Study time and activities C3: Family background C4: School characteristics	C1: Grade, age, sex, immigrant, speaks foreign language, occupational aspiration C2: time spent at school, homework is assigned C3: parental education, parental occupational status, wealth, books present at home C4: % of girls, % repeating grade, class size, student-teacher ratio, teacher qualifications, access to computers/internet, shortage of teachers by subject PISA survey (2006)	> Explained component Varies considerably across countries; e.g. in maths: 0% (Germany), 51% (Korea), 100% (Japan) > Categories (specific results reported for Korea and Japan) Korea: explained component attributable to differences in C4 (63%-83%) Japan: explained component attributable to differences in C2 and C4 (> 85%) Remarks: Analysis on Korea and Japan illustrates how to identify factors that contribute most to the gap. If the observed gap is mainly due to the "unexplained" component, public policy needs to focus on broader and underlying economic, social and cultural differences.
9	Burger (2011) Zambia	Test scores in reading Urban vs. rural schools	Regressions on contemporaneous predictors OB twofold; reference coefficients from rural schools Components: endowments ("explained"), coefficients ("unexplained")	None	Predictors: household asset index, parental education, pupil-teacher ratio, proportion of students who spoke English Southern Africa Consortium for Monitoring Educational Quality (SACMEQ II)	> Endowments ("resources"): 55% > Coefficients ("returns to resources"): 45% Policy implication: Resource investment will not have the required impact unless the efficiency gap is also addressed.

	(A) Author(s), year and country	(B) Outcome and groups	(C) Empirical specification, decomposition strategy and components	(D) Categories proposed for observed influences	(E) Observed influences and data sources	(F) Results, remarks and policy implications
10	Beltran and Seinfeld (2012) Peru	Test scores in reading and maths Urban vs. rural students	Regressions on contemporaneous predictors. OB threefold; rural students as reference group. Components: endowments, returns, interaction.	None	Predictors: preschool attendance, district classified as poor, parental education, mother tongue, teacher qualifications, school infrastructure, access to internet at school, % repeating grade at school, private- public school, class time (minutes) National Student Evaluation (2010) and School Census (2010)	> Endowments: 36% (reading); 22% (maths) > Coefficients/returns: 14.4% (reading); 4.8% (maths) > Interaction: 49.6% (reading); 73.2% (maths) Policy implication: Adequate provision of resources has to be complemented with quality assurance and effective use.

Appendix 3.1

From the cumulative to the value added specification with age-dependent parameters

Consider the following linear version of the skill formation technology with age-dependent parameters.

$$A_{i2} = HI'_{i2}\gamma_1^2 + HI'_{i1}\gamma_2^2 + SI_{i2}\phi_1^2 + h_{i2}\varphi_1^2 + h_{i1}\varphi_2^2 + f'_i\lambda_{(2)} + \mu_{i0}\beta_{(2)} \quad (i)$$

Superscript 2 in all the parameters in (i) implies they are specific to period 2. This implies that the effect of a certain input on skill not only depends on the time elapsed since the application of the input but also on the specific period (or age) in which it was applied. Based on this, period 1 skill can be expressed as:

$$A_{i1} = HI'_{i1}\gamma_1^1 + h_{i1}\varphi_1^1 + f'_i\lambda^1 + \mu_{i0}\beta^1 \quad (ii)$$

We can subtract ρA_{i1} from (i) to obtain:

$$\begin{aligned} A_{i2} - \rho A_{i1} &= HI'_{i2}\gamma_1^2 + HI'_{i1}(\gamma_2^2 - \rho\gamma_1^1) + SI_{i2}\phi_1^2 + h_{i2}\varphi_1^2 + h_{i1}(\varphi_2^2 - \rho\varphi_1^1) \\ &\quad + f'_i(\lambda_{(2)} - \rho\lambda^1) + \mu_{i0}(\beta_{(2)} - \rho\beta^1) \end{aligned} \quad (iii)$$

If the effects of inputs decay at a rate ρ , the effects of period 1 inputs on period 2 skill equal ρ times the effect of period 1 inputs on period 1 skill. This means: $\gamma_2^2 = \rho\gamma_1^1$, $\varphi_2^2 = \rho\varphi_1^1$, $\lambda_{(2)} = \lambda^2 + \rho\lambda^1$ and $\beta_{(2)} = \beta^2 + \rho\beta^1$. This, in turn, implies (iii) can be written as:

$$A_{i2} = \rho A_{i1} + HI'_{i2}\gamma_1^2 + SI_{i2}\phi_1^2 + h_{i2}\varphi_1^2 + f'_i\lambda^2 + \mu_{i0}\beta^2 \quad (iv)$$

The main text presents the case for age-invariant parameters with essentially the same results (see equations (3.24), (3.25) and (3.26)). The only difference is that parameters accompanying home, school and health inputs, as well as predetermined direct influences and innate ability, are no longer age or period-specific.

Appendix 3.2

Coefficient estimates for the eight specifications involved

Table 3A

Estimations with the reduced sample (including all variables and excluding school inputs)

VARIABLES	“Full information” estimations		Excluding school inputs (reduced sample)	
	Hybrid-CU	Hybrid-VA	Hybrid-CU	Hybrid-VA
Real expenditure in child (learning materials and entertainment; round 2)	0.076 (0.090)	-0.022 (0.083)	0.180 (0.106)	0.038 (0.087)
Mother had antenatal visits during pregnancy (yes = 1)	0.132** (0.051)	0.121** (0.052)	0.136** (0.050)	0.119** (0.053)
Maternal response to child cry was affectionate (yes = 1)	0.077 (0.063)	0.025 (0.070)	0.101 (0.062)	0.035 (0.065)
Child attended formal preschool (yes = 1)	0.049 (0.094)	0.003 (0.080)	0.050 (0.121)	0.009 (0.113)
Household has books and child is encouraged to read (yes = 1)	0.210*** (0.057)	0.235*** (0.068)	0.207*** (0.052)	0.240*** (0.062)
Household has a computer (yes = 1)	0.087 (0.059)	0.064 (0.055)	0.133* (0.075)	0.099 (0.072)
Real expenditure in child (learning materials and entertainment; round 3)	0.062 (0.062)	0.036 (0.064)	0.077 (0.058)	0.041 (0.064)
Child receives help from parents when doing homework (yes = 1)	0.007 (0.093)	-0.041 (0.091)	0.027 (0.085)	-0.023 (0.085)
Hours in a typical day the child spends playing	-0.003 (0.021)	-0.012 (0.026)	0.032* (0.017)	0.015 (0.022)
Hours in a typical day the child spends sleeping	-0.032 (0.036)	-0.044 (0.038)	-0.011 (0.032)	-0.028 (0.035)
Hours in a typical day the child spends studying	0.045 (0.045)	0.024 (0.041)	0.085* (0.046)	0.047 (0.046)
Child is stunted (yes = 1; round 2)	-0.048 (0.085)	-0.002 (0.084)	-0.079 (0.101)	-0.025 (0.094)
Child is stunted (yes = 1; round 3)	-0.212 (0.123)	-0.184 (0.133)	-0.226 (0.134)	-0.198 (0.143)
Hours in a typical day the child spends at school	-0.075 (0.051)	-0.070 (0.053)	-0.042 (0.056)	-0.043 (0.054)
Years of schooling (basic education)	0.342*** (0.081)	0.253*** (0.075)	0.326*** (0.077)	0.227*** (0.060)
CLIM: absence of problems in class (score 12-48)	0.009** (0.004)	0.011** (0.004)	--	--
INF: school has basic services (yes = 1)	0.183** (0.069)	0.044 (0.064)	--	--
ACT: average curricular coverage (% of topics covered in depth)	0.521 (0.308)	0.388 (0.267)	--	--
ORG: teacher absenteeism (%)	-0.780* (0.380)	-0.761** (0.317)	--	--
ORG: school has a psychologist (yes = 1)	0.194** (0.079)	0.203** (0.087)	--	--
ORG: school is “multigrade” (yes = 1)	-0.292** (0.120)	-0.308*** (0.098)	--	--

VARIABLES	“Full information” estimations		Excluding school inputs (reduced sample)	
	Hybrid-CU	Hybrid-VA	Hybrid-CU	Hybrid-VA
TEA: more than 50% of teachers graduated from university (yes = 1)	0.090* (0.049)	0.012 (0.046)	--	--
Child’s caregiver has higher education (yes = 1)	0.135* (0.066)	0.017 (0.056)	0.169** (0.075)	0.025 (0.058)
Caregiver’s age	0.013** (0.005)	0.008 (0.005)	0.012** (0.005)	0.007 (0.005)
Child is male (yes = 1)	0.028 (0.096)	0.034 (0.083)	0.056 (0.094)	0.062 (0.078)
Child’s mother tongue is Spanish (yes = 1)	0.341** (0.144)	0.389** (0.152)	0.390** (0.145)	0.439** (0.170)
Child’s age in months	0.014 (0.011)	0.004 (0.009)	0.018 (0.011)	0.007 (0.010)
Child lives in urban area (yes = 1)	0.120 (0.113)	0.028 (0.096)	0.395*** (0.076)	0.202** (0.091)
Average household total income	0.018 (0.024)	0.010 (0.021)	0.033 (0.032)	0.021 (0.027)
Average household size	0.013 (0.026)	0.013 (0.026)	0.006 (0.028)	0.006 (0.028)
Proportion of male siblings	-0.032 (0.175)	-0.100 (0.173)	-0.137 (0.146)	-0.202 (0.138)
Child birth order	-0.097*** (0.026)	-0.085** (0.031)	-0.101*** (0.025)	-0.088*** (0.029)
Caregiver aspiration for child is university education (yes = 1)	0.058 (0.054)	0.026 (0.059)	0.112 (0.072)	0.065 (0.070)
Standardized raw PPVT score (round 2)		0.341*** (0.052)		0.363*** (0.046)
Constant	-1.362 (1.176)	0.465 (1.117)	-1.894 (1.152)	0.102 (1.172)
Observations	487	487	487	487
R-squared	0.557	0.608	0.510	0.573

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3B

Estimations with the complete sample excluding school inputs and with the reduced sample including school fixed effects

VARIABLES	Excluding school inputs (complete sample)		School fixed effects (reduced sample)	
	Hybrid-CU	Hybrid-VA	Hybrid-CU	Hybrid-VA
Real expenditure in child (learning materials and entertainment; round 2)	0.122*** (0.036)	0.034 (0.028)	-0.022 (0.099)	-0.103 (0.096)
Mother had antenatal visits during pregnancy (yes = 1)	0.193*** (0.044)	0.139*** (0.039)	0.212*** (0.062)	0.205** (0.071)
Maternal response to child cry was affectionate (yes = 1)	0.042 (0.043)	0.012 (0.046)	0.038 (0.062)	0.012 (0.056)
Child attended formal preschool (yes = 1)	0.059 (0.058)	0.016 (0.044)	0.070 (0.169)	0.043 (0.147)
Household has books and child is encouraged to read (yes = 1)	0.183*** (0.048)	0.194*** (0.048)	0.283** (0.105)	0.272** (0.110)
Household has a computer (yes = 1)	0.130*** (0.040)	0.082** (0.034)	0.074 (0.090)	0.061 (0.083)
Real expenditure in child (learning materials and entertainment; round 3)	0.059 (0.044)	0.029 (0.032)	0.145 (0.098)	0.098 (0.108)
Child receives help from parents when doing homework (yes = 1)	0.076 (0.045)	0.061 (0.044)	0.020 (0.149)	-0.023 (0.136)
Hours in a typical day the child spends playing	0.037* (0.021)	0.021 (0.018)	-0.037 (0.032)	-0.051 (0.031)
Hours in a typical day the child spends sleeping	-0.019 (0.024)	-0.036 (0.024)	-0.038 (0.032)	-0.054 (0.038)
Hours in a typical day the child spends studying	0.065*** (0.022)	0.017 (0.018)	0.025 (0.069)	0.015 (0.063)
Child is stunted (yes = 1; round 2)	-0.098 (0.066)	-0.024 (0.050)	-0.054 (0.089)	-0.022 (0.097)
Child is stunted (yes = 1; round 3)	-0.125** (0.052)	-0.111** (0.053)	-0.227* (0.122)	-0.190 (0.152)
Hours in a typical day the child spends at school	0.020 (0.037)	0.005 (0.034)	-0.048 (0.060)	-0.057 (0.058)
Years of schooling (basic education)	0.254*** (0.044)	0.134*** (0.034)	0.309** (0.022)	0.286** (0.103)
CLIM: absence of problems in class (score 12-48)	--	--	--	--
INF: school has basic services (yes = 1)	--	--	--	--
ACT: average curricular coverage (% of topics covered in depth)	--	--	--	--
ORG: teacher absenteeism (%)	--	--	--	--
ORG: school has a psychologist (yes = 1)	--	--	--	--
ORG: school is "multigrade" (yes = 1)	--	--	--	--
TEA: more than 50% of teachers graduated from university (yes = 1)	--	--	--	--

VARIABLES	Excluding school inputs (complete sample)		School fixed effects (reduced sample)	
	Hybrid-CU	Hybrid-VA	Hybrid-CU	Hybrid-VA
Child's caregiver has higher education (yes = 1)	0.194*** (0.058)	0.054 (0.048)	0.071 (0.101)	-0.059 (0.088)
Caregiver's age	0.009** (0.004)	0.006 (0.004)	0.018* (0.010)	0.014 (0.010)
Child is male (yes = 1)	0.068 (0.059)	0.060 (0.038)	0.021 (0.114)	0.056 (0.102)
Child's mother tongue is Spanish (yes = 1)	0.278** (0.105)	0.367*** (0.100)	0.343 (0.264)	0.316 (0.268)
Child's age in months	0.012 (0.007)	0.008 (0.006)	0.016 (0.009)	-0.002 (0.009)
Child lives in urban area (yes = 1)	0.453*** (0.085)	0.246*** (0.083)	-0.101 (0.152)	-0.131 (0.150)
Average household total income	0.048** (0.018)	0.027* (0.014)	0.001 (0.040)	-0.000 (0.032)
Average household size	0.009 (0.014)	0.008 (0.013)	0.039 (0.040)	0.042 (0.042)
Proportion of male siblings	-0.012 (0.069)	-0.029 (0.052)	-0.149 (0.219)	-0.218 (0.225)
Child birth order	-0.076*** (0.020)	-0.051** (0.021)	-0.105** (0.042)	-0.097* (0.048)
Caregiver aspiration for child is university education (yes = 1)	0.216*** (0.034)	0.147*** (0.033)	0.122 (0.081)	0.073 (0.083)
Standardized raw PPVT score (round 2)		0.399*** (0.030)		0.352*** (0.064)
Constant	-1.667* (0.807)	-0.249 (0.813)	-0.614 (0.697)	1.851* (0.895)
School fixed effects	No	No	Yes	Yes
Observations	1,561	1,561	487	487
R-squared	0.473	0.557	0.698	0.730

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3C**Testing interactions between school quality variables and past cognitive achievement**

VARIABLES	Hybrid-VA	Hybrid-VA with interactions
Real expenditure in child (learning materials and entertainment; round 2)	-0.022 (0.083)	-0.023 (0.082)
Mother had antenatal visits during pregnancy (yes = 1)	0.121** (0.052)	0.110* (0.061)
Maternal response to child cry was affectionate (yes = 1)	0.025 (0.070)	0.034 (0.066)
Child attended formal preschool (yes = 1)	0.003 (0.080)	0.005 (0.079)
Household has books and child is encouraged to read (yes = 1)	0.235*** (0.068)	0.231*** (0.064)
Household has a computer (yes = 1)	0.064 (0.055)	0.072 (0.060)
Real expenditure in child (learning materials and entertainment; round 3)	0.036 (0.064)	0.044 (0.066)
Child receives help from parents when doing homework (yes = 1)	-0.041 (0.091)	-0.049 (0.086)
Hours in a typical day the child spends playing	-0.012 (0.026)	-0.008 (0.030)
Hours in a typical day the child spends sleeping	-0.044 (0.038)	-0.044 (0.039)
Hours in a typical day the child spends studying	0.024 (0.041)	0.022 (0.040)
Child is stunted (yes = 1; round 2)	-0.002 (0.084)	-0.006 (0.089)
Child is stunted (yes = 1; round 3)	-0.184 (0.133)	-0.172 (0.125)
Hours in a typical day the child spends at school	-0.070 (0.053)	-0.066 (0.051)
Years of schooling (basic education)	0.253*** (0.075)	0.266*** (0.075)
S1: absence of problems in class (score 12-48)	0.011** (0.004)	0.012** (0.004)
S2: school has basic services (yes = 1)	0.044 (0.064)	0.053 (0.077)
S3: average curricular coverage (% of topics covered in depth)	0.388 (0.267)	0.326 (0.227)
S4: teacher absenteeism (%)	-0.761** (0.317)	0.040 (0.446)
S5: school has a psychologist (yes = 1)	0.203** (0.087)	0.201** (0.092)
S6: school is “multigrade” (yes = 1)	-0.308*** (0.098)	-0.267* (0.124)
S7: more than 50% of teachers graduated from university (yes = 1)	0.012 (0.046)	-0.000 (0.038)
Child’s caregiver has higher education (yes = 1)	0.017 (0.056)	0.031 (0.066)
Caregiver’s age	0.008 (0.005)	0.008 (0.005)

VARIABLES	Hybrid-VA	Hybrid-VA with interactions
Child is male (yes = 1)	0.034 (0.083)	0.049 (0.083)
Child's mother tongue is Spanish (yes = 1)	0.389** (0.152)	0.413** (0.155)
Child's age in months	0.004 (0.009)	0.001 (0.009)
Child lives in urban area (yes = 1)	0.028 (0.096)	0.000 (0.092)
Average household total income	0.010 (0.021)	0.009 (0.024)
Average household size	0.013 (0.026)	0.013 (0.027)
Proportion of male siblings	-0.100 (0.173)	-0.141 (0.167)
Child birth order	-0.085** (0.031)	-0.087** (0.030)
Caregiver aspiration for child is university education (yes = 1)	0.026 (0.059)	0.026 (0.062)
Standardized raw PPVT score (round 2)	0.341*** (0.052)	0.510* (0.248)
Standardized raw PPVT score*S1	--	-0.000 (0.005)
Standardized raw PPVT score*S2	--	0.033 (0.066)
Standardized raw PPVT score*S3	--	-0.403 (0.248)
Standardized raw PPVT score*S4	--	1.107 (0.694)
Standardized raw PPVT score*S5	--	0.031 (0.148)
Standardized raw PPVT score*S6	--	0.094 (0.136)
Standardized raw PPVT score*S7	--	0.035 (0.056)
Constant	0.465 (1.117)	0.702 (1.071)
Observations	487	487
R-squared	0.608	0.615
Joint significance of interactions		
F-stat	--	1.16
p-value		0.33

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Chapter 4

Effects of Early Childhood Development Interventions on Parental Behaviour: Evidence from a Home-Visiting Programme in Peru

Abstract

Home-visiting interventions aim at improving child outcomes by shifting caregivers' parenting skills so they can offer better stimulation opportunities to their children. Despite this explicit objective, there is little evidence from the developing world regarding the effects of the home-visiting model on the provision of stimulation by caregivers, and no evidence from programmes working at scale. The analysis presented in this chapter seeks to contribute to the literature on ECD interventions by providing causal evidence about the effects of a scaled-up home-visiting programme on caregiver behaviour. I use the cluster-randomised design of the home-visiting component of programme *Cuna Mas* in Peru to collect data on parenting practices from 20 control and 20 treatment rural districts. A novel feature of the survey instruments is that they can account for stimulation opportunities offered by the caregiver in addition to the activities carried out during the home visit, which means they can better capture changes in caregiver behaviour. Results reveal that: (i) treatment effects on the quality of the home environment are positive, statistically significant ($p < 0.01$) and have a size ($d = 0.5$) comparable to that found for other interventions of much more smaller scale conducted in the developing world; and (ii) the intervention has caused an increase in the variety of play activities offered to the child by the caregiver in addition to those occurring during home visits ($d = 0.3$; $p < 0.01$). The latter provides strong evidence of a shift in caregiver behaviour. This shift is important because it increases the likelihood that an enriched environment will be offered to the child beyond the duration of the programme.

4.1 Introduction and motivation

The lack of adequate early stimulation at home has captured attention as a major constraint preventing children in socioeconomically disadvantaged families from reaching better developmental outcomes (Heckman, 2006; Walker et al., 2007). It is well established that parent-child interactions are crucial in shaping child development during their first years of life (Huberman and Mendelsohn, 2012) and evidence suggests that parental engagement in

educational play activities with their children is a crucial input for child development (Fiorini and Keane, 2014).

Early childhood development (ECD) interventions aim at improving one or more dimensions of child development⁷⁷ by enhancing the inputs received by the child during her first years of life. This can be attempted directly (for example, by providing food supplementation or offering an enriched play environment at a day care centre) or indirectly (for example, by offering caregivers the skills and materials for them to provide an enriched environment to their children).

The home-visiting model constitutes a prominent example of an ECD intervention that combines a direct and an indirect effect. It offers direct stimulation and materials during home visits carried out by a trained paraprofessional or community member, usually once per week and for a period of 1 to 2 years. Importantly, home visits also seek to improve caregivers' parenting skills for them to be able to offer better stimulation opportunities and enriched interactions to their children. This type of interventions are therefore classified among those focused on parent-child interactions (Huberman and Mendelsohn, 2012).

The literature offers strong evidence to support that ECD interventions can deliver improvements in child outcomes (Nores and Barnett, 2010). Four systematic reviews of ECD interventions in the developing world also provide strong evidence to support that the home-visiting model can cause this type of improvements (Walker (2011), Engle et al. (2011), Baker-Henningham and Lopez-Boo (2010), Engle et al. (2007)). There is, however, much less evidence regarding the effects of these interventions on the quality of the home environment and, in particular, on parental behaviour conducive to the provision of better stimulation opportunities to the child. In addition, and to the best of my knowledge, no

⁷⁷ The early childhood development literature distinguishes at least three developmental outcomes or skills: sensorimotor, cognitive and socio-emotional (Boyden and Dercon, 2012; Grantham-McGregor et al., 2007).

evaluation so far has assessed the impacts of a scaled-up home-visiting programme and there have been very few attempts to use evaluation results to try to understand the mechanisms behind caregivers' behavioural change.

In light of the above, the analysis presented in this and the following chapter seeks to contribute to the literature on ECD in two ways. First, it provides causal evidence about the effects of a scaled-up home-visiting intervention on parental behaviour. And, second, it explores the constraints that limit caregivers' behavioural change and proposes a way of thinking about the mechanisms behind it. For this, the analysis is organized around three research questions: (i) can a home-visiting ECD intervention working at scale in rural Peru deliver an improvement in the quality of the home environment?; (ii) can a home-visiting ECD intervention working at scale in rural Peru change parental behaviour so as to increase the amount of stimulation offered to the child; and (iii) what constraints faced by the caregiver limit the effect of this intervention on caregiver engagement in educational play activities with the child?

This chapter focuses on the first two research questions and is organized as follows. Section 4.2 provides an overview of the available evidence regarding the effect of home-visiting ECD interventions on parental behaviour and introduces some key concepts. This should serve to justify and locate my research questions within the relevant literature. Section 4.3 describes the intervention considered for this study and explains the research design. The effects of the intervention on the quality of the home environment and parental behaviour are discussed in Section 4.4. Finally, Section 4.5 presents some concluding remarks related to this part of the analysis.

The third research question is addressed in Chapter 5. In particular, I explore the heterogeneity of treatment effects on caregiver behaviour and relate this to the heterogeneity

found in parental beliefs regarding the importance of parent-child interactions for child development. I also propose a simple model to interpret these findings and offer a way of thinking about the mechanisms behind caregivers' behavioural change.

4.2 Home-visiting ECD interventions, the home environment and parental behaviour: evidence from the developing world

A review of the ECD literature on home-visiting interventions in the developing world reveals three features that this research seeks to address: (i) scant evidence about the effect of these interventions on parental behaviour, together with the fact that its measurement can be problematic because instruments will typically pick-up the effect of the home visit which does not imply a change in parental behaviour; (ii) lack of evidence from programmes or interventions working at scale; and (iii) little attention paid to the mechanisms through which these interventions affect parental behaviour and a strong emphasis on the transfer of parenting skills to caregivers ("parent education"). This emphasis carries the implicit assumption that caregivers' lack of knowledge about certain activities that can be performed with their children and of the ability to engage in them are the major binding constraint preventing parents from providing more stimulation opportunities to their children.

The literature on ECD interventions provides strong evidence to support the claim that the home-visiting model can deliver positive results in terms of child outcomes in the developing world. Four systematic reviews ((Walker (2011), Engle et al. (2011), Baker-Henningham and Lopez-Boo (2010), Engle et al. (2007)) identify 16 different studies that have evaluated the effects of home visits on early child development outcomes (ages between 0 and 3) and all of them report benefits in at least one domain of childhood development. In addition, results from the recent evaluation of a home-visiting intervention in Colombia also reveal positive results on child outcomes (Attanasio et al., 2012; Attanasio et al., 2013).

Home-visiting interventions have the explicit objective of improving caregivers' parenting practices by demonstrating to them how to engage in play activities with their children and how to be more sensitive and responsive during daily interactions. This seeks to ensure an improvement in the quality of the home environment additional to that directly offered during the home visit, and which might even extend beyond the duration of the intervention.

In the developing world, the provision of better stimulation opportunities at home constitutes an important component of home-visiting interventions in response to the significant cognitive skill setbacks found among children belonging to disadvantaged families (Grantham-McGregor et al., 2007). In fact, the literature identifies the lack of adequate early stimulation at home as one of the main risk factors preventing poor children from reaching better developmental outcomes (Walker et al., 2007).

Despite this explicit interest in parenting practices and stimulation opportunities, there is much less evidence regarding the effects of the home-visiting model on the home environment or the provision of stimulation by parents. In only 9 out of the 17 studies mentioned above, the evaluation comprised measurement of the effects on the quality of the home environment. Moreover, while aggregate positive results were found in 7 of these 9 studies, only 3 report positive effects in terms of stimulation opportunities offered to the child by the caregiver.⁷⁸

Instruments commonly employed to account for the effect of ECD interventions on the home environment include culturally adapted versions of items from the Home Observation for Measurement of the Environment (HOME) inventory (Caldwell and Bradley, 1984) or from the Family Care Indicators (FCI) early childhood development module (UNICEF,

⁷⁸ These three studies are: Attanasio et al. (2013) (significant effect in the variety of play activities and play material offered to the child), Powell et al. (2004) (significant effect on mothers' childrearing practices), and Walker et al. (2004) (significant effect on maternal involvement with the child). See Appendix 4.1 for the complete list of studies.

2011). Based on the structure of these instruments and the information usually collected in the studies referred above, we can conceive the quality of the home environment as a concept that involves at least one of the following two aspects: (i) the amount of stimulation opportunities offered to the child (which can include child's engagement in particular activities, her access to play material and the organization of the child's personal space); and (ii) the degree of caregivers' sensitivity and responsiveness (which refer to caregivers' ability to understand her child's needs and cues and her ability to respond to these needs and cues, respectively). Caregiver sensitivity and responsiveness are usually evaluated in terms of the verbal and affective interactions between the caregiver and the child (Totsika and Sylva, 2004).

This way of understanding the quality of the home environment is closely related to the concept of parenting practices. In fact, Hoff et al. (2002) refer to the HOME inventory when discussing how are parenting practices defined and measured. These authors acknowledge there is little consensus on how to conceptualize "parenting practices" and decided to focus on caregiver-child interactions. They organized their analysis considering three aspects: (i) verbal interaction; (ii) direct control practices (how controlling, restrictive and punitive parents are with their children); and (iii) managerial control (which refers to the experiences and physical environments parents provide to their children).

The concept of behavioural change I propose for this analysis implies a shift in parenting practices that leads to an improvement in the quality of the home environment additional to that offered during home visits, and which is sustained, at least, for the duration of the intervention. Based on this, and despite being closely related, it is better to keep the concepts of home environments and parenting practices (or parental behaviour) separate when studying the effects of a home-visiting intervention. This is especially relevant if

follow-up data is collected while the intervention is still in progress (as in this study) or at the same time as it is being phased out. This is because the intervention offers a direct improvement in the quality of the home environment through home visits that can involve the caregiver, but which do not necessarily imply that the caregiver is providing an additional and sustained improvement.

From the above it follows that an increase in the amount of stimulation offered to the child at home that is provoked by an on-going home-visiting intervention does not necessarily imply there has been a behavioural change among caregivers. The HOME inventory and the FCI collect information about the stimulation opportunities offered to the child by observation (for example, confirming the presence of play material) and by direct report of the primary caregiver (regarding her or other adults' involvement in particular activities with the child). During the intervention or while it is being phased out, changes can be observed or reported in these two aspects only because of what has happened during the home visits. The timing of the follow-up survey, the source of the play material, and the frequency and people involved in the interactions are, therefore, important elements that have to be considered when analysing the effects of home-visiting interventions on parental behaviour. It is not clear from the three studies referred above whether these elements were considered when measuring the effects of the intervention.⁷⁹

It should also be noted that none of the nine studies that analysed the effects of the home-visiting model on the home environment used results from an intervention working at scale. Most of the studies can be classified as efficacy trials that worked with a limited number of children and families (samples sizes below 150) and delivered treatment under ideal or

⁷⁹ In two of the three studies, follow-up data was collected at the same time as the intervention was being phased out (Attanasio et al., 2013; Powell et al., 2004). In one study, information regarding the home environment was collected midway through the intervention (Walker et al., 2004). None of the studies explicitly explain whether the provision of play material and caregivers' engagement in play activities was additional to that occurring during home visits.

highly controlled conditions (e.g. working with a small group of trained professionals or paraprofessionals as home-visitors). The intervention more akin to a scaled-up programme is that evaluated in Attanasio et al. (2013).⁸⁰ Generating evidence from scaled-up interventions is important for policy because it provides a more realistic appraisal of what can be accomplished with a large target population. This is especially relevant in the developing world because targeted groups (e.g. families living below a poverty line) are usually large and the home-visiting model is highly dependent on the personal skills of home-visitors, which means that the quality of delivery can be particularly sensitive to scale.

Finally, and despite the obvious emphasis that home-visiting and other PCI-focused interventions have on parental behaviour, little is explicitly said in the studies referenced above about the mechanisms through which they affect parenting practices. Most studies that address the issue of mechanisms focus on the channels linking the intervention with child development.⁸¹ This emphasis is surely relevant in terms of policy outcomes, but it tends to overlook a process which is far from mechanical and which determines, at least in part, the success of these interventions.

Despite the literature's lack of explicit attention to the mechanisms behind caregivers' behavioural change, a review of its emphasis in terms of the characteristics of successful interventions can provide some clues regarding the implicit assumptions being made about

⁸⁰ The intervention was linked to the administrative capacity and community networks of the Colombian conditional transfer programme (*Familias en Accion*) and home-visits were carried out by community members for whom special training and material were prepared. This intervention also targeted a large sample of children (1,429).

⁸¹ In Baker-Henningham and Lopez Boo (2010), for example, the authors highlighted: (i) how early stimulation can ameliorate the negative effects associated with living in disadvantaged circumstances; (ii) the benefits on mothers' parenting behaviour; (iii) the benefits to maternal mental health; and (iv) improvements in children's ability to take advantage of other educational opportunities (increased school readiness). Nothing was explicitly said, however, about how changes in mothers' parenting behaviour are produced. In a more recent study, Attanasio et al. (2015) acknowledge that these interventions can modify parental investment decisions by making them aware of the importance of early investments. Their analysis, however, did not explicitly consider this or other potential channels for caregivers' behavioural change. In fact, these authors focused on the mechanisms that explain the shift in child outcomes and used reduced form parental investment equations with parameters that were allowed to vary with treatment status.

the process. What we find in this regard is a strong emphasis on the transfer of skills to parents.

For example, in one very influential Lancet review on ECD, the authors conclude that parenting interventions have larger effects when they include systematic curricula and training opportunities for home visitors and parents, as well as “active strategies to show and promote caregiving behaviours – e.g. practice, role play, or coaching to improve parent-child interactions” (Engle et al., 2011; p. 1343). Other authors also stress the importance of designing programmes to be participatory and interactive with parents (Huberman and Mendelsohn, 2012).

Parenting skills or ability, understood in this case as knowledge of particular play activities to engage with children and of how to implement them, are surely necessary for parents to be able to offer more stimulation opportunities to their children. However, a strong emphasis on this sole component implicitly assumes that the lack of these skills is the only binding constraint preventing parents in poor families from having a more active role in influencing their children’s early environment. This is a strong assumption to make if resources, in general, are scarce, as is the case among the populations these programmes are meant to serve.

In this regard, Huberman and Mendelsohn (2012) noted that PCI-focused programmes have produced larger benefits with relatively better educated families. These authors cite U. Bronfenbrenner (1974) on the role of living conditions, who notes that “in many homes, the conditions of life are so harsh that, so long as they persist, the parent has neither the will nor the capacity to participate in educational activities with the child” (Bronfenbrenner, 1974; p. 36). Of course, “disadvantaged”, “at risk” or “harsh conditions” can describe a rather wide range of circumstances when referring to household resources. Further analysis of the

constraints limiting caregivers' behavioural change is therefore required, and this analysis will involve an exploration of the possible mechanisms explaining this phenomenon.

4.3 The intervention and research design

4.3.1 *Cuna Mas* home-visiting programme

The rural component of the *Cuna Mas* ECD programme in Peru (called *Servicio de Acompañamiento a Familias*) is based on the Jamaica Home Visiting Model developed by Sally Grantham-McGregor and Christine Powell (Huberman and Mendelsohn, 2012). It comprises the provision of 1-hour home visits every week for a period of one and a half years.

Home visits are carried out by specially trained community members or *facilitadoras*. Each *facilitadora* is in charge of approximately 10 families. They follow a structured curriculum (PNCM, 2013) to: (i) show caregivers how to interact in a more sensitive and responsive way with their children during basic care routines; (ii) demonstrate to caregivers how to engage in age-appropriate play activities with their children; and (iii) show caregivers book sharing techniques. The intervention also comprises the delivery of children's books, illustrations, and simple toys such as puzzles.

The intervention was focused at the district level. Districts selected are those: (i) where monetary poverty incidence is above 50%; (ii) that fall within the scope of intervention of the conditional cash transfer programme *Juntos*; (iii) that exhibit a rate of chronic malnutrition for children aged 0-5 above 30%; and (iv) where more than 50% of the population lives in a rural community. These criteria yielded a total of 531 districts comprising 31,828 rural communities. Children between 0 and 3 years of age living in these communities represent the target population of the intervention, which will be progressively

expanded seeking full coverage by year 2016. The first scaling-up effort of the programme occurred between May and June 2013 and, by mid-2015, it was serving approximately 52,000 families across 14 regions.⁸² The data for this study were collected between February and March 2014.

4.3.2 Evaluation design

The intervention was scaled-up in a way that allowed for an evaluation with experimental design. In particular, 60 districts were randomly allocated to each of the following groups: home visits (“treatment 1”), home visits and group sessions (“treatment 2”) and no intervention (control).⁸³ The two largest rural communities in terms of target population were selected within each district and, from each community, a total of 15 children-caregiver dyads were randomly selected to participate in a baseline survey collected between April and May 2013 (MIDIS, 2013). Follow-up data collection for the evaluation of *Cuna Mas*’s effects on child development was programmed for 2015.

This study is based on the cluster-randomised design described above, and has the objective of estimating the effectiveness of the intervention in changing the quality of the home environment and caregivers’ behaviour. Following power calculations, it was decided to randomly select 20 districts from the control group and 20 districts from the “treatment 2” group already created. Within each district, in turn, we planned to randomly select 15 caregivers from the baseline sample. This yielded a total planned sample of 600 caregivers, 300 from the control group and 300 from the “treatment 2” group (henceforth referred to as “treatment group”). This sample size and distribution ensured a minimum detectable size of 0.26 standard deviations for a statistical power of 80%. This was judged as a conservative

⁸² According to the information in the 2014 and 2015 budgets, the average cost per family is around US\$ 450 per year.

⁸³ The map presented in Appendix 4.2 shows the distribution of districts. Group sessions were discontinued in early 2014 following very low coverage due to implementation problems.

minimum effect size considering the values found in the studies referred above (where effect sizes ranged between 0.32 and 0.65).

The intervention was not launched at the same time in all communities belonging to the treatment group because of administrative delays in regional offices. In fact, by the time fieldwork was implemented, 45 out of the 120 treatment communities had not started receiving the intervention. The study had a limited budget⁸⁴ and the Programme wanted the largest possible sample coverage of those communities where they had started operations. Accordingly, the Programme required that the treatment sample should include only those caregivers living in communities where the intervention had been launched. Combined with the condition of surveying an average of 15 caregivers per district and the fact that baseline data were available from two communities per district, this requirement implied that 30 treatment districts had to be excluded from the analysis. In fact, excluded districts were those that had less than 15 baseline caregivers living in a community where the intervention had started by the time fieldwork was implemented.

After these exclusions, 20 districts were randomly selected from the remaining control and treatment groups. In each selected district, in turn, 15 caregivers were randomly selected for interview while the rest, if any, were randomly allocated to a replacement list (Table 4A in Appendix 4.2 presents the final list of districts and regions involved in this study).

In principle, the exclusions described above could have introduced bias in the results. For example, if exclusions affected more distant districts, one could expect behavioural differences between control and included treatment caregivers driven by disparities in wealth or educational attainment. There is, however, no evidence of this. In fact, there are

⁸⁴ Budget limitations also implied that 3 regions (out of 11) had to be excluded from the analysis in order to prevent excessive geographical dispersion of the sample. This implied losing 14 districts from the original control group and 8 districts from the original treatment group.

no significant differences in terms of baseline household, child and caregiver characteristics and, more importantly, in terms of outcome measures, between: (i) control and treatment districts available after the exclusion of regions; (ii) included and excluded districts within the treatment group after consideration of there being a minimum of 15 baseline caregivers living in a community where the intervention had started by the time fieldwork was implemented; (iii) included districts within the treatment group and districts in the control group; and (iv) excluded districts within the treatment group and districts in the control group (see Appendix 4.3).

Absence of statistically significant differences for all these permutations provides strong evidence in favour of the comparability of the control and treatment groups considered. In particular, if treatment districts not affected by exclusions are comparable to control districts, the same should apply to all the remaining subgroups. In addition, the fact that all the groups share the same pre-treatment outcome measures in terms of caregiver behaviour provides stronger evidence to support the absence of bias. This is because the presence of unobservable confounders affecting treatment estimates would likely manifest by producing unbalanced pre-treatment behaviours.

It is also worth recalling that the main reason why some treatment communities had not been offered treatment by the time fieldwork was implemented was because of administrative delays in regional offices, a reason which is, in principle, unrelated to the behaviours I aim to analyse. In addition, the exclusion rule was influenced by the requirements of power calculations and this meant that the criteria for not being part of the treatment sample at the caregiver level amounted to inhabiting a district where the intervention had not yet started in, at least, one of its two communities. In most of the cases, excluded districts had only one community that had not started to receive the intervention.

This reduced the likelihood of exclusion being correlated with a characteristic shared by all district members.

Departure from a pure random design means that the absence of unobservable confounders is not assured. It is, however, a reasonable assumption based on the arguably exogenous nature of the phenomenon driving exclusion (living in a district where the intervention has not yet started in, at least, one of its communities) and on the fact that districts included in the treatment group share the same pre-treatment characteristics and outcome measures as those excluded, and as those belonging to the control group.

An analysis of baseline information for the control and treatment caregivers actually surveyed in the districts that were finally selected is presented in the next section. This confirms that both groups shared balanced behaviours and characteristics prior to the intervention.

Administrative records available for the treatment sample contain information on the number of visits provided and the date of the first and last visit delivered until December 2013. These records reveal that caregivers in the treatment sample had received an average of 25 visits until December 2013, a figure consistent with the fact that nearly 6 months had passed since the programme was scaled up (between May and June 2013). A refusal rate can be estimated considering those caregivers living in communities that have started to receive treatment and who register zero visits. This figure is close to 8%. A dropout rate can be approximated considering those caregivers whose last recorded visit occurred before December 2013. According to this criterion, 3% of caregivers were no longer receiving treatment by that date.

4.3.3 Survey instruments

A questionnaire was prepared to measure the quality of the home environment as well as caregivers' beliefs regarding the importance of early parent-child interactions for child development (see Appendix 4.4).

Items included in this questionnaire allow one to characterize the quality of the home environment on the basis of four aspects: (i) interaction and play activities (variety of caregiver-child interactions during basic care and variety of play activities offered to the child); (ii) responsiveness and control practices (communicative and affective caregiver-child interactions and how the caregiver disciplines the child); (iii) play material and home conditioning (variety of play materials present at home and the conditions of the child's play area); and (iv) personal care and hygiene (overall child and caregiver appearance in terms of hygiene). Figure 4B in Appendix 4.4 provides more detail on these specific components and indicates whether the information was collected through caregivers' report or enumerators' observation during the interview. This information was used to build a quality of the home environment index (with values ranging 0-1) given by the simple average of the scaled scores obtained in each of the four components described above.⁸⁵

It should be noted that the abovementioned components comprise the two basic aspects of the concept of quality of home environment discussed in section 4.2; i.e. (i) the amount of stimulation opportunities offered to the child; and (ii) the degree of caregiver sensitivity and responsiveness. In addition, the information organized under these four components encompasses all the aspects considered by the six subscales of the HOME infant-toddler

⁸⁵ Figure 4B in Appendix 4.4 also indicates the raw scores that can be obtained in each component. Note that maximum raw scores are not the same across components. Scaling with respect to the maximum score within each component ensures components have equal weights in the aggregate index. As shown below, results are robust to using the raw score obtained after adding up all items (i.e. without scaling within each component). Building the aggregate index this way ensures all items receive the same weight.

questionnaire as well as most of the information collected by the FCI early childhood development module. In fact, specific items used to measure caregivers' responsiveness and control practices and the variety of play material are adapted versions of the items contained in the *Emotional and verbal responsivity*, *Avoidance of restriction and punishment*, and *Provision of appropriate play materials* subscales of the HOME inventory. Items used to measure the variety of play activities are based on the FCI.

In terms of characterizing the quality of the home environment, the questionnaire employed for this study includes two novel elements with respect to the HOME inventory and the FCI: (i) specific items aimed at capturing the variety of caregiver-child interactions during feeding and basic care routines; and (ii) a modified version of the play activity table used in the FCI that allows one to identify both the frequency of the interaction and the person interacting with the child. As will be discussed later, both elements are important to account for the effects of the intervention on the stimulation opportunities offered to the child in addition to those provided during the home visit.

Caregiver beliefs regarding the importance of parent-child interactions for child development were captured through an eight-item scale (see items 550 and 551 in the survey questionnaire presented in Appendix 4.4). During the interview, caregivers were asked about the ideal educational attainment and occupation they would like their child to achieve and were then asked to what extent they agree or disagree with eight statements about the relation between early parent-child interaction and these goals. Measurement of these beliefs is also one of the novel features of this analysis and will serve, in the next chapter, to provide evidence regarding the mechanisms behind caregivers' behavioural change.

The questionnaire was piloted in three rural districts with similar socioeconomic characteristics to those involved in the evaluation. Interviews were conducted in the language commonly used by caregivers (around 35% spoke Quechua). Informed consent was obtained from all participants without revealing that the study was related to *Cuna Mas*. In fact, this information (as well as treatment status) was also concealed from enumerators to prevent them from making reference to the programme during the interview.

4.3.4 Balance between treatment and control groups at baseline

Examination of control and treatment groups' observable characteristics prior to treatment can provide evidence to support the claim that there are no confounders influencing the results obtained after comparing outcome measures between these two groups. As already noted, evidence is even stronger if these groups share the same pre-treatment outcome measures as the presence of unobservable confounders affecting treatment estimates would likely manifest by producing unbalanced pre-treatment behaviours.

Table 4.1 summarizes mean values in the control and treatment group for pre-treatment outcome measures related to the quality of the home environment and several household, caregiver and child characteristics relevant for caregivers' child-rearing practices. Measures were built using baseline data collected between April and May 2013. A quality of the home environment index was built organizing baseline data under the four components described above. The specific information entering each component, however, is not exactly the same as the one considered in this study because the baseline survey was not as rich as our instrument in terms of measuring caregiver behaviour (for example, it did not collect information on the variety of caregiver-child interactions during basic care). The household wealth index was built combining information on dwelling characteristics, access to basic services and availability of durable goods, as described in Escobal et al. (2003).

Table 4.1**Outcome variables and household, caregiver and child characteristics at baseline**

	Control	Treatment	Difference
<i>Outcome variables</i>			
Quality of home environment index (0-1)	0.599	0.604	0.005 (0.02)
Interaction and play activities (0-1)	0.425	0.434	0.009 (0.028)
Responsiveness and control practices (0-1)	0.708	0.717	0.009 (0.024)
Play material and home conditioning (0-1)	0.348	0.363	0.015 (0.03)
Personal care and hygiene (0-1)	0.914	0.902	-0.012 (0.026)
<i>Household, caregiver and child characteristics</i>			
Household wealth index (0-1)	0.481	0.460	-0.021 (0.036)
Household members	5.425	5.395	-0.030 (0.175)
Caregiver's age	29.334	28.221	-1.113 (0.732)
Caregiver's educational attainment (years)	6.214	6.827	0.612 (0.593)
Caregiver is married (yes = 1)	0.306	0.221	-0.085 (0.057)
Caregiver worked outside household last week (yes = 1)	0.619	0.769	0.150* (0.078)
Number of children under caregiver's care	1.034	1.044	0.010 (0.014)
Caregiver's first language is Spanish (yes = 1)	0.650	0.629	-0.021 (0.131)
Caregiver is the mother (yes = 1)	0.975	0.973	-0.002 (0.015)
Child's age (months)	12.863	13.129	0.267 (0.595)
Child is male (yes = 1)	0.531	0.476	-0.055 (0.044)
Number of observations	320	294	

Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

In general, we can say control and treatment characteristics and outcome measures were balanced prior to the intervention. It should also be noted from Table 4.1 that the final

sample comprises 614 observations, 320 (52%) belonging to the control group and 294 (48%) belonging to the treatment group. As planned, data was collected from 20 control districts and 20 treatment districts with an average of 15.4 observations per district. Accordingly, all standard errors are corrected to consider that the data is arranged in 40 clusters.

4.4 Treatment effects on the quality of the home environment and parental behaviour

4.4.1 Estimation of treatment effects

Treatment effects were calculated using OLS estimates of parameters β and β^C in equations (4.1) and (4.2) given below.

$$y_i = \alpha + \beta T_i + \mu_i \quad (4.1)$$

$$y_i = \alpha^C + \beta^C T_i + x_i' \gamma + \varepsilon_i \quad (4.2)$$

In these equations, y_i is the outcome measure under analysis, T_i is the treatment status of the caregiver ($T_i = 1$ if she belongs to the treatment group, $T_i = 0$ if she belongs to the control group), and x_i is a vector containing pre-treatment characteristics and outcome measures.

Balanced pre-treatment characteristics and outcome measures should translate into similar estimates for β and β^C . Estimates provided by equation (4.2), however, are usually preferred because they can produce more precise estimates of the causal effect of interest (Angrist and Pischke, 2009). Note that the fact that covariates are not correlated with the treatment implies that their exclusion will not bias the estimate of β . However, this does not imply they have no explanatory power for y_i and, therefore, their inclusion can reduce the residual variance.

At this point is worth noticing that being part of the treatment group implies that the caregiver has been offered treatment but not necessarily that she has accepted it. Thus, in principle, the parameter estimated in (4.1) and (4.2) corresponds to an intention to treat analysis. Given that the refusal rate is low (around 8% according to the estimate discussed above), however, one can claim that this intent-to-treat effect is equivalent to the effect of the treatment on the treated. That is, the effect of the intervention itself and not the effect of being assigned to treatment (Duflo et al., 2007).

Recall that the Programme required that the treatment sample should include only those caregivers living in communities where the intervention had been launched. This implies that the intent-to-treat estimate provided in this analysis does not include the effect of delays in starting the intervention. This should not be considered a problem since the empirical goal is to estimate the effectiveness of home visits (the intervention itself) in changing the quality of the home environment and caregivers' behaviour.⁸⁶

4.4.2 Effects on the quality of the home environment

Table 4.2 presents treatment effects on the quality of the home environment index and the scaled scores of each component (complete regression results are reported in Appendix 4.5). Effects are reported in terms of the simple difference between control and treatment group means and as the difference controlled for pre-treatment outcome values and household, caregiver and child characteristics considered in Table 4.1. Results for the aggregate index are reported considering the average over all components. This implies losing observations from caregivers not responding all items or for whom all characteristics could not be

⁸⁶ If data from treatment communities where the intervention had not yet started were available, one could use the original treatment assignment as an instrument for the actual treatment status to obtain the intent-to-treat estimate provided here. In the previous section, I discussed and presented empirical evidence to support the claim that the rule determining the exclusion of districts from the treatment group is exogenous. This claim is comparable to postulating that the results provided by the instrumental variable strategy just explained would have been equivalent to the results presented here.

observed during the interview.⁸⁷ It is worth mentioning that pre-treatment characteristics and outcome measures are also balanced for this subsample of caregivers. In addition, results shown in the first row of Table 4.2 are robust to building the aggregate index averaging only over those components with complete information and to considering the aggregate score obtained after adding up all items.⁸⁸

Table 4.2
Treatment effects on the quality of the home environment

	Obs.	Mean control group	Mean treatment group	Difference	Controlled difference ^(a)	d ^(b)
<i>Aggregate index (0-1)</i>						
Average over all components	486 (C=238; T=248)	0.610	0.671	0.062*** (0.019)	0.059*** (0.013)	0.534
<i>Index components (0-1)</i>						
1. Interaction and play activities	614 (C=320; T=294)	0.363	0.438	0.074*** (0.019)	0.068*** (0.013)	0.382
2. Responsiveness and control practices	548 (C=283; T=265)	0.644	0.707	0.063*** (0.021)	0.062*** (0.021)	0.424
3. Play material and home conditioning	537 (C=264; T=273)	0.490	0.592	0.101*** (0.03)	0.098*** (0.023)	0.511
4. Personal care and hygiene	558 (C=290; T=268)	0.939	0.945	0.006 (0.016)	0.01 (0.011)	0.101

Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

(a) Controlled difference is the estimated effect including pre-treatment controls.

(b) Effect size (d) calculated as the standardized controlled difference.

Overall, evidence shows that the treatment has a positive and significant effect on the quality of the home environment. Its effect size (standardized treatment effect) is around 0.5. This size can be regarded as moderate (Cohen, 1992) and lies between the values found

⁸⁷ Most missing observations are due to the fact that some characteristics of the environment could not be observed because the caregiver was not willing to show the play material or particular areas of the house to the enumerator. Rejection was low (around 10%) but its effect gets compounded when building an aggregate score.

⁸⁸ Averaging over those components with complete information allows one to use the complete sample of 614 caregivers. Results are as follows: mean control group = 0.597; mean treatment group = 0.657; controlled difference = 0.057 (0.014); d = 0.466. Considering the aggregate score obtained after just adding up all items ensures all items are given the same weight. Results are as follows: mean control group = 0.595; mean treatment group = 0.657; controlled difference = 0.059 (0.013); d = 0.524.

in the literature for effective trials based on the home-visiting model: $d = 0.65$ in Jamaica (Powell et al., 2004); $d = 0.53-0.54$ in Colombia (Attanasio et al., 2013); $d = 0.49$ in Jamaica (Grantham-McGregor et al., 1991); $d = 0.37$ in Jamaica (Walker et al., 2004); $d = 0.32$ in India (Bentley et al., 2010).

It should be noted that the abovementioned interventions had between 1 and 2 years of duration. This is significantly above the average exposure of the *Cuna Mas* treatment group at the time this data was collected (around 6 months or an average of 25 visits). It is reasonable to assume, however, that longer or more intense exposures are more related to larger treatment effects in terms of child outcomes than in terms of caregiver behaviour.⁸⁹

Table 4.2 also reveals that the improvement in the quality of the home environment is explained by positive treatment effects on every component except in “personal care and hygiene”. It should be noted that the score obtained in this component by the control group is already high (0.94 out of 1) leaving little space for an improvement.

4.4.3 Effects on caregiver behaviour

Achieving behavioural change among caregivers is important because it implies an enriched environment is being offered to the child in a continuous way and not only during home visits. A change in caregivers’ behaviour also increases the likelihood of this improvement being sustained even after the intervention is phased out.

As already noted, an increase in stimulation opportunities offered to the child (either in terms of play activities or play materials) does not imply there has been a change in parental behaviour if results stem from follow up data that has been collected while home visits are

⁸⁹ In Powell and Grantham-McGregor (1989) the evaluation design allowed one to compare the effects of having biweekly vs. monthly visits for a period of 2 years. Larger effects on child outcomes were found for biweekly visits. Unfortunately, no measurements of caregiver behaviour were taken.

still being carried out. This is the case of this particular evaluation. In addition, it is also prudent to refrain from considering increases in access to play material when talking about changes in caregiver behaviour if the intervention has directly provided materials and it is not possible to account for their origin from survey data.

To measure the effect of the intervention on caregiver behaviour we will focus on caregiver-child interactions happening in addition to the home visit. The instruments used for this analysis allow one to do this by looking at two subcomponents of the quality of the home environment index: (i) the number of types (or variety) of interactions during basic care and play; and (ii) the number of types (or variety) of play activities offered to the child. The first one is built considering interactions that take place independently of the home visit (e.g. verbal communication that happens during feeding, bathing and clothes change routines). As shown in the first row of Table 4.3, there is a positive and significant effect in this outcome variable.

The second outcome variable is of particular interest for this analysis because it is related to activities with a direct and explicit stimulation purpose. This indicator is built considering the number of types of activities (out of 7 possible types) carried out by the caregiver with the child at least twice per week. Conditioning on a minimum frequency of 2 implies that the variety is measured excluding those types of activities carried out only once per week. This is important because it allows one to approximate the increase in the variety of play activities that take place in addition to what happens during the home visit.⁹⁰ Results

⁹⁰ Excluding those activities carried out with frequency 1 when building the indicator implies excluding the effect of treatment on the number of activities carried out once per week. This provides a good approximation to the effect of treatment on the number of activities offered to the child that take place in addition to what happens during the home visit, based on the notion that the main effect of the intervention is to allow caregivers to start doing activities they were not offering to the child. If these activities are performed twice per week or more, it means they are being carried out in addition to what takes place during the home visit. Also note that the results (presented in the second row of Table 4.3) are robust to measuring treatment effects after subtracting 1 from the frequencies reported by treatment caregivers for activities that are likely to happen

obtained for this indicator confirm that the treatment has produced an increase in the variety of play activities offered to the child by the caregiver, over and above the stimulation provided during the home visit (see the second row in Table 4.3).

Table 4.3
Treatment effects on caregivers' interaction with child and on the variety of play activities offered by the caregiver

	Mean control group	Mean treatment group	Difference	Controlled difference^(a)	d^(b)
Number of types of interactions during basic care (0-12)	5.069	5.813	0.744*** (0.235)	0.747*** (0.208)	0.301
Number of types of play activities carried out at least twice last week by the caregiver (0-7)	1.838	2.500	0.663*** (0.214)	0.537*** (0.135)	0.310

Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

(a) Controlled difference is the estimated effect including pre-treatment controls.

(b) Effect size (d) calculated as the standardized controlled difference.

Another benefit of collecting information about the frequency and person involved in play activities is that one can shed light on how the change in the child's environment is occurring. One can think of the nature of the environment to which the child is exposed during play as something that depends, at least, on the type of activity being carried out and the people involved with the child during the activity. In principle, changes can occur in either characteristic.

When we asked about child engagement in particular play activities and its weekly frequency we also recorded who was the person usually sharing the activity with the child and if there was somebody else involved. This allows one to decompose the total effect on

during a home visit (one obtains 0.49 for the controlled difference; $p < 0.01$). This is a conservative estimate because not all the activities are necessarily carried out during each home visit.

the variety of play activities carried out with a minimum frequency of f ($\Delta V(f)$) in the following way:

$$\begin{aligned}\Delta V(f) &= \Delta V(f|\textit{without caregiver involvement}) \\ &\quad + \Delta V(f|\textit{with caregiver involvement})\end{aligned}\tag{4.3}$$

Notice that caregiver involvement implies that the caregiver was either identified first as the person usually sharing the activity with the child or was identified as the one accompanying the interaction between the first person and the child. This implies that the effect on the variety of play activities carried out with the child with caregiver involvement can be further decomposed following:

$$\begin{aligned}\Delta V(f|\textit{with caregiver involvement}) &= \Delta V(f|\textit{caregiver identified as first}) \\ &\quad + \Delta V(f|\textit{caregiver identified as accompanying})\end{aligned}\tag{4.4}$$

Table 4.4 presents the results of these two decompositions. As explained below, they reveal that increased stimulation offered to the child (additional to that occurring during the home visit) is occurring through greater caregiver involvement in play activities rather than through the introduction of more types of play activities.

Table 4.4

Effects on the variety of play activities offered to the child (number of types of play activities)

Person involved	Minimum frequency	
	1	2
(a) Anyone ($\Delta V(f)$) [(b) + (c)]	0.73*** (0.113)	0.274** (0.116)
(b) Without caregiver involvement	-0.272* (0.150)	-0.508*** (0.137)
(c) With caregiver involvement [(d) + (e)]	1.001*** (0.15)	0.782*** (0.137)
(d) Caregiver identified first as the person usually sharing the activity	0.666*** (0.159)	0.537*** (0.135)
(e) Caregiver identified as the person accompanying the interaction between the first person and the child	0.340*** (0.079)	0.249*** (0.061)

Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The first row in Table 4.4 shows there is only a small effect of around 0.27 on the number of types of activities offered to the child that are additional to those that take place during the home visit (i.e. with a minimum frequency of 2). Moreover, if we compare columns 1 and 2 for indicator (a) we would realise that most of the effect on the variety of play activities is concentrated on frequency 1, something that could cast doubt on the effectiveness of the intervention beyond the home visit.⁹¹

This result, however, masks the fact that there is greater caregiver involvement in play activities, something that could not have been identified in absence of information regarding the people involved in these activities. In fact, the effect of the intervention on the variety of play activities carried out with a minimum frequency of 2 raises from 0.27 to 0.78 when we condition it to have caregivers' involvement (see indicator (c) in Table 4.4). According

⁹¹ The effect on the variety of play activities carried out with a minimum frequency of f ($\Delta V(f)$) can be decomposed as the sum of the effects on each particular frequency starting in f . Formally: $\Delta V(f) = \Delta V_f + \Delta V_{f+1} + \dots + \Delta V_7$. Following the results presented in the first row of Table 4.4, this implies that the effect on the variety of play activities carried out only once per week is given by $\Delta V_1 = \Delta V(1) - \Delta V(2) = 0.73 - 0.27 = 0.46$.

to equation (4.3), for this to be possible, the variety of play activities offered without caregiver participation must be falling by $0.27 - 0.78 = -0.51$, as confirmed by the results reported for indicator (b) in Table 4.4.⁹²

These results imply that the nature of the environment to which the child is exposed during play, in addition to the home visit, is being transformed mainly through the engagement of their caregivers (i.e. a shift in the people involved) and less because of the introduction of more types of activities.⁹³

It is also worth noting that the indicator proposed to measure the variety of play activities offered to the child by the caregiver when accounting for changes in caregiver behaviour (reported in Table 4.3 above) corresponds to that conditioned on the caregiver being identified first as the person usually sharing the activity with the child (indicator (d) in Table 4.4). This is to ensure that the caregiver has a leading role in the activity carried out with the child.

The effect on the number of types of play activities is equal to the sum of the effects on the proportion of caregivers who engage in each of the activities considered. This allows one to analyse if there are certain activities that account for the majority of the overall effect on the variety of play offered to the child.

Table 4.5 allows this analysis by showing the effect on the proportion of caregivers who engage in each of the 7 types of play activities considered. Three types of activities explain most of the effect on the variety of play offered to the child by the caregiver: telling stories

⁹² This result should not be taken to imply that a substitution effect is in place. Greater caregiver involvement does not necessarily mean that other household members are no longer taking part of the activity.

⁹³ Increasing the variety of play activities performed with the caregiver is not the only way to enhance child outcomes. Additional interaction with other family members or increases in the frequency of activities can also have a positive effect. This analysis focuses on the variety of play activities performed with the caregiver based on the objectives of the Programme and the causal evidence linking caregiver-child interactions with child development. An analysis of other mechanisms (e.g. participation of other family members) constitutes an interesting opportunity for future research.

to the child (10 percentage point increase), singing songs to or with the child (12 percentage point increase), and playing with the child and her toys (16 percentage point increase). No significant effect has been found in activities such as reading or looking at books, drawing or painting, and taking the child outside the house. The proportion of caregivers who carry out this last activity, however, is already relatively large in the control group (around 60%).

Table 4.5
Treatment effects on each type of play activity (proportion of caregivers)

Play activities		Mean control group	Mean treatment group	Controlled difference
<i>Play activities</i>				
1	Read books, look at pictures from a book with (NAME)	0.200	0.252	0.032 (0.031)
2	Told stories to (NAME)	0.100	0.208	0.102*** (0.03)
3	Sang songs to or with (NAME), including lullabies	0.303	0.439	0.122** (0.048)
4	Took (NAME) outside the house	0.606	0.605	-0.007 (0.045)
5	Played with (NAME) with his/her toys	0.166	0.350	0.161*** (0.04)
6	Drew, painted or scribbled with (NAME)	0.169	0.245	0.054 (0.035)
7	Played with (NAME) to name objects, colours or numbers	0.294	0.401	0.073* (0.038)
			Sum^(a) d	0.537*** 0.310

(a) The sum of effects corresponds to the overall effect on the variety of play activities carried out by the caregiver at least twice per week.

Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

4.5 Concluding remarks

Despite the explicit interest that home-visiting ECD interventions implemented in the developing world have in enhancing stimulation opportunities for children, there is scant

evidence regarding their effect on the home environment and, especially, on caregivers' engagement in educational play activities with their children. Moreover, there is a lack of evidence from interventions working at scale while the potential populations to be targeted are large and the home-visiting model can be especially sensitive to size as it is heavily dependent on the proficiency of home-visitors.

The analysis presented in this chapter aimed at contributing to the ECD literature by providing causal evidence about the effects of a scaled-up home-visiting intervention on parental behaviour. For this, two research questions were addressed: (i) can a home-visiting ECD intervention working at scale in rural Peru deliver an improvement in the quality of the home environment?; (ii) can a home-visiting ECD intervention working at scale in rural Peru change parental behaviour and increase the amount of stimulation offered to the child?

The answers to these research questions follow directly from the treatment effects found in the analysis. In particular, estimated treatment effects of the *Cuna Mas* home-visiting programme on the quality of the home environment are positive, statistically significant ($p < 0.01$) and have a size ($d = 0.5$) comparable to that found for other interventions of much smaller scale and efficacy trials conducted in the developing world (Attanasio et al., 2013; Bentley et al., 2010; Grantham-McGregor et al., 1991; Powell et al., 2004; Walker et al., 2004).

Treatment effects on caregiver behaviour related to the provision of stimulation opportunities for the child are also positive and statistically significant. In fact, it was possible to identify an increase in the variety of play activities offered to the child by the caregiver in addition to those occurring during home visits ($d = 0.3$; $p < 0.01$). Positive evidence of a shift in caregiver behaviour is important because this increases the likelihood that an enriched environment will be offered to the child beyond the duration of the

programme. The process of skill formation is cumulative and early investments can increase the productivity of future inputs (Cunha and Heckman, 2008; Cunha, Heckman, and Schennach, 2010). Because of this, a permanent improvement in the environment being offered to the child can help sustain and even boost the effect of the intervention on child development.

Appendix 4.1

Studies that have evaluated interventions with a home-visiting component

	Study	Country	Duration	Child outcomes		Quality of home environment	
				Measured effects	Positive results	Measured effects	Positive results
1	Cooper et al. (2002)	South Africa	0.5 years	Height and weight	Yes	Sensitivity and responsiveness	Yes
2	Cooper et al. (2009)	South Africa	5 months	Socio emotional development	Yes	Sensitivity and responsiveness	Yes
3	Gardner et al. (2003)	Jamaica	8 weeks	Cognitive development	Yes	No	NA
4	Powell and Grantham-McGregor (1989)	Jamaica	1.5 years	Cognitive development	Yes	No	NA
5	Powell (2004)	Jamaica	1 year	Cognitive development	Yes	Stimulation opportunities	No
6	Magwaza and Edwards (1991)	South Africa	10 weeks	Cognitive development	Yes	No	NA
7	Powell et al. (2004) Baker-Henningham et al. (2005)	Jamaica	1 year	Cognitive development	Yes	Stimulation opportunities	Yes
8	Walker et al. (2004) Walker et al. (2010)	Jamaica	2 years	Cognitive and socio emotional development	Yes	Stimulation opportunities	Yes
9	Grantham-McGregor et al. (1983) Grantham-McGregor et al. (1987) Grantham-McGregor et al. (1989) Grantham-McGregor et al. (1994)	Jamaica	3 years	Cognitive development	Yes	Stimulation opportunities	No
10	Grantham-McGregor et al. (1991)	Jamaica	2 years	Cognitive and socio emotional development	Yes	Aggregate HOME score	Yes
11	Bentley et al. (2010)	India	NA	Cognitive development and motor skills	Yes	Aggregate HOME score	Yes
12	Janssens and Rosenberg (2011)	St Lucia	1 year	Cognitive development	Yes	No	NA

	Study	Country	Duration	Child outcomes		Quality of home environment	
				Measured effects	Positive results	Measured effects	Positive results
13	Attanasio et al. (2013)	Colombia	1.5 years	Cognitive development	Yes	Stimulation opportunities	Yes
14	Eickmann et al. (2003)	Brazil	5 months	Cognitive and motor development	Yes	No	NA
15	Hamadani et al. (2006)	Bangladesh	1 year	Cognitive development	Yes	No	NA
16	Nahar et al. (2009)	Bangladesh	6 months	Cognitive and motor development	Yes	No	NA
17	Nair et al. (2009)	South India	1 year	Cognitive and motor development	Yes	No	NA

Appendix 4.2

Districts and regions involved in the study

Figure 4A
Districts selected for the evaluation exercise of the rural component of programme *Cuna Mas*




Table 4A
Districts and regions included in the final sample

	Control group			Treatment group		
	Region	Province	District	Region	Province	District
1	Apurimac	Andahuaylas	Santa Maria de Chicmo	Apurimac	Antabamba	Juan Espinoza Medrano
2	Cajamarca	Cajabamba	Cachachi	Apurimac	Aymaraes	Caraybamba
3	Cajamarca	Cajamarca	Cospan	Cajamarca	Cutervo	Santo Domingo de la Capilla
4	Cajamarca	Cajamarca	Jesús	Cajamarca	Hualgayoc	Bambamarca
5	Cajamarca	Jaen	Colasay	Cajamarca	Jaen	Chontali
6	Cajamarca	San Ignacio	San Jose de Lourdes	Cajamarca	Jaen	San Jose Del Alto
7	Cusco	Paucartambo	Huancarani	Cajamarca	San Ignacio	La Coipa
8	Huancavelica	Acobamba	Anta	Cajamarca	Santa Cruz	Catache
9	Huancavelica	Huancavelica	Huachocolpa	Cusco	Canas	Yanaoca
10	Huancavelica	Tayacaja	Akraquia	Cusco	Chumbivilcas	Capacmarca
11	Huanuco	Huamalties	Monzón	Cusco	Paucartambo	Paucartambo
12	Huanuco	Huánuco	Santa Maria del Valle	Huanuco	Huacaybamba	Huacaybamba
13	Huanuco	Leoncio Prado	Hermilio Valdizan	Huanuco	Huacaybamba	Pinra
14	Huanuco	Pachitea	Umari	Huanuco	Huamalties	Miraflores
15	Junin	Huancayo	Pucara	Huanuco	Leoncio Prado	Mariano Damaso Beraun
16	Junin	Tarma	Tapo	Junin	Chupaca	Yanacancha
17	Junin	Yauli	Huay-huay	Junin	Concepcion	Chambara
18	La Libertad	Sánchez Carrion	Sanagoran	La Libertad	Santiago De Chuco	Quiruvilca
19	Puno	Carabaya	Ollachea	Piura	Huancabamba	Sondorillo
20	Puno	San Antonio De Putina	Quilcapuncu	Puno	Carabaya	Usicayos

Appendix 4.3

Quality of home environment, household, caregiver and child characteristics in the districts considered for the evaluation

	Control (a)	Treatment (b)	Diff. (a) – (b)	Treatment Included (c)	Treatment Excluded (d)	Diff. (c) – (d)	Diff. (a) – (c)	Diff. (a) – (d)
Quality of home environment index (0-1)	0.583	0.595	-0.012 (0.012)	0.608	0.585	0.022 (0.019)	-0.024 (0.015)	-0.002 (0.014)
Household wealth index (0-1)	0.442	0.443	-0.001 (0.019)	0.459	0.432	0.027 (0.026)	-0.017 (0.024)	0.01 (0.023)
Household members	5.495	5.384	0.111 (0.1)	5.406	5.368	0.039 (0.153)	0.089 (0.125)	0.127 (0.11)
Caregiver's age	28.365	28.144	0.221 (0.31)	28.118	28.163	-0.045 (0.386)	0.247 (0.421)	0.202 (0.367)
Caregiver's educational attainment (years)	6.250	6.636	-0.386 (0.33)	6.989	6.376	0.613 (0.479)	-0.739* (0.408)	-0.126 (0.377)
Caregiver is married (yes = 1)	0.253	0.221	0.031 (0.03)	0.214	0.227	-0.013 (0.039)	0.039 (0.038)	0.026 (0.037)
Caregiver worked outside household last week (yes = 1)	0.635	0.689	-0.054 (0.042)	0.757	0.639	0.118** (0.053)	-0.121** (0.052)	-0.004 (0.05)
Number of children under caregiver's care	1.023	1.031	-0.007 (0.006)	1.036	1.027	0.009 (0.008)	-0.012* (0.007)	-0.003 (0.006)
Caregiver's first language is Spanish (yes = 1)	0.584	0.639	-0.055 (0.081)	0.592	0.673	-0.081 (0.115)	-0.008 (0.104)	-0.089 (0.093)
Caregiver is the mother (yes = 1)	0.970	0.979	-0.009 (0.007)	0.983	0.975	0.008 (0.008)	-0.014 (0.009)	-0.006 (0.009)
Child's age (months)	12.738	13.011	-0.273 (0.275)	13.280	12.814	0.467 (0.358)	-0.542 (0.363)	-0.076 (0.32)
Child is male (yes = 1)	0.490	0.511	-0.021 (0.021)	0.506	0.515	-0.008 (0.031)	-0.016 (0.027)	-0.025 (0.023)
Number of observations	46	52	--	22	30	--	--	--

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix 4.4

Survey questionnaire and characterization of the quality of the home environment

A Parenting practices		
> Let us talk about the activities you carry out with (NAME)		
1. Responsive breastfeeding and feeding		
501	Enumerator, indicate: ¿Is the child 6 months old or more?	1 Yes >> skip to question 510
		2 No >> continue with question 502
502	Are you currently breastfeeding or bottle-feeding (NAME)?	1 Yes, breastfeeding >> continue with question 503
		2 Yes, bottle-feeding >> continue with question 503
		3 No >> skip to question 513
503	What do you do with (NAME) while you breastfeed him/her? Enumerator > do not read the options, mark all that apply	1 Looks at the child
		2 Talks to the child, sings or reads to him/her
		3 Caresses child
		4 Other (specify) _____
		5 Does nothing with child
504	Do you perform any other activity (housework, for example) while you breastfeed (NAME)?	1 Yes, while breastfeeding the child, she also... Specify _____
		2 No, no other activity, she only breastfeeds the child
505	Between yesterday afternoon and today's morning, has (NAME) had any liquid or food other than breast milk?	1 Yes Specify _____ >> skip to section 2: Interaction with child
		2 No >> skip to section 2: Interaction with child
510	Are you currently breastfeeding or bottle-feeding (NAME)?	1 Yes, breastfeeding >> continue with question 511
		2 Yes, bottle-feeding >> continue with question 511
		3 No >> skip to question 513
511	What do you do with (NAME) while you breastfeed him/her? Enumerator > do not read the options, mark all that apply	1 Looks at the child
		2 Talks to the child, sings or reads to him/her
		3 Caresses child
		4 Other (specify) _____
		5 Does nothing with child
512	Do you carry out any other activity (housework, for example) while you breastfeed (NAME)?	1 Yes, while breastfeeding the child, she also... Specify _____
		2 No, no other activity
513	Is (NAME) currently eating any solid foods?	1 Yes >> continue with question 514
		2 No >> skip to section 2: Interaction with child
514	Who does usually feed (NAME)?	1 Primary caregiver
		2 Activity is not carried out by primary caregiver > indicate code _____ 2. (NAME)'s mother; 3. (NAME)'s father; 4. (NAME)'s brothers or sisters; 5. (NAME)'s grandparents; 6. Other
		3 The child eats by him/herself
515	What do you do with (NAME) while he/she eats? Enumerator > do not read the options, mark all that apply	1 Talks to the child
		2 Plays or sings with child
		3 Encourages child to participate in the activity; teaches child how to use a spoon or fork
		4 Allows child to participate in the activity; allows him/her to use a spoon or fork
		5 Other (specify) _____
		6 Does nothing with child
516	Do you carry out any other activity (housework, for example) while (NAME) eats?	1 Yes, while the child eats, she... Specify _____
		2 No, no other activity


517	Usually, (NAME) eats: Enumerator, choose all the answers that apply	1	Breakfast					
		2	Mid-morning snack					
		3	Lunch					
		4	Mid-afternoon snack					
		5	Dinner					
		6	Other (specify) _____					
2. Interaction with child								
			Yes	No	With whom?	Anyone else?	No. of days	
520	During the last 7 days, has anyone shared with (NAME) the following activities? Enumerator, consider the following codes to identify the person > 1. Primary caregiver <i>If it is someone different than the primary caregiver, use the following codes ></i> 2. (NAME)'s mother 3. (NAME)'s father 4. (NAME)'s brothers or sisters 5. (NAME)'s grandparents 6. Other	1	Read books, look at pictures from a book with (NAME)	1	2			
		2	Told stories to (NAME)	1	2			
		3	Sang songs to or with (NAME), including lullabies	1	2			
		4	Took (NAME) outside the house	1	2			
		5	Played with (NAME) with his/her toys	1	2			
		6	Drew, painted or scribbled with (NAME)	1	2			
		7	Played with (NAME) to name objects, colours or numbers	1	2			
521	What do you do with (NAME) while he/she plays? Enumerator > do not read the options, mark all that apply	1	Talks to the child					
		2	Encourages child to participate in play					
		3	Allows child to play freely					
		4	Introduces challenging situations during play					
		5	Other (specify) _____					
		6	Does nothing with child					
522	Would you like to spend more time reading books, singing or playing with (NAME)?	1	Yes					
		2	No >> skip to question 524					
523	Why can't you spend more time reading books, singing or playing with (NAME)? Enumerator > do not read the options, mark all that apply	1	I have other things to do / I am busy / I need to work					
		2	The child plays with his/her brothers and sisters					
		3	The child gets tired, falls asleep					
		4	I would not know how to play with him/her					
		5	Other (specify) _____					
524	Who does usually bath (NAME)?	1	Primary caregiver					
		2	Activity is not carried out by primary caregiver > indicate code _____ 2. (NAME)'s mother; 3. (NAME)'s father; 4. (NAME)'s brothers or sisters; 5. (NAME)'s grandparents; 6. Other					
		3	The child washes him/herself					
525	What do you do with (NAME) during his/her bath? Enumerator > do not read the options, mark all that apply	1	Talks to the child					
		2	Plays or sings with child					
		3	Encourages child to participate in the activity; teaches child how to wash him/herself					
		4	Allows child to participate in the activity; allows him/her to wash him/herself					
		5	Other (specify) _____					
		6	Does nothing with child					
526	Do you carry out any other activity (housework, for example) during (NAME)'s bath?	1	Yes Specify _____					
		2	No, no other activity					
527	Who does usually change (NAME)'s clothes?	1	Primary caregiver					
		2	Activity is not carried out by primary caregiver > indicate code _____ 2. (NAME)'s mother; 3. (NAME)'s father; 4. (NAME)'s brothers or sisters; 5. (NAME)'s grandparents; 6. Other					
		3	The child changes his/her own clothes					

528	What do you do with (NAME) while changing his/her clothes? Enumerator > do not read the options, mark all that apply	1	Talks to the child
		2	Plays or sings with child
		3	Encourages child to participate in the activity; teaches child how to change his/her clothes
		4	Allows child to participate in the activity; allows him/her to change his/her clothes
		5	Other (specify) _____
		6	Does nothing with child
529	Do you carry out any other activity (housework, for example) while (NAME) eats?	1	Yes Specify _____
		2	No, no other activity
3. Home conditioning			
> I would like to know more about the places of the house and objects that (NAME) uses			
530	Could you show me the place where (NAME) usually plays?	1	Yes
		2	No >> skip to question 532
531	Enumerator: observe and choose all characteristics that apply	1	The place is clean (not littered)
		2	It is well illuminated
		3	It is well ventilated
		4	Has a blanket or similar to prevent direct contact with the floor
		5	It is not near fire, objects that might injure the child, stairways, or the road
		6	None of the above
532	Could you show me the toys or objects that (NAME) usually uses to play?	1	Yes
		2	No >> skip to question 534
		3	Has no toys / uses no toys >> skip to question 536
533	Enumerator: observe and choose all types of toys that are present	1	Children books
		2	Materials to draw or paint
		3	Cuddly or role playing toys
		4	Push or pull toys
		5	Building blocks or puzzles
		6	Toys to play music
		7	None of the above
534	Could you show me the place where (NAME)'s toys are stored	1	Yes
		2	No >> skip to question 536
535	Enumerator: observe and choose all characteristics that apply	1	Has a box, bag or similar to store the toys
		2	Has no specific place, toys are around the house
536	Could you show me the place where you usually eat?	1	Yes
		2	No >> skip to question 538
537	Enumerator, observe: is it a place where the family can share a meal?	1	Yes
		2	No
538a	Could you show me the place where you usually wash or clean yourselves?	1	Yes
		2	No >> skip to the next section (question 540)
539a	Enumerator: observe and choose all characteristics that apply Do they have a water tap or bucket with clean water?	1	Yes
		2	No
538b	Could you show me the utensils you use for personal hygiene?	1	Yes
		2	No >> skip to the next section (question 540)
539b	Enumerator: observe and choose all characteristics that apply	1	They have a soap and a soapdish (an element that prevents contact of soap with dirt)
		2	They have a towel or cloth
		3	None of the above

4. Supervision and discipline management			
540	Sometimes, adults taking care of children have to leave the house to go shopping, wash clothes, etc. and have to leave the children. During the last 7 days...		¿How many times was (NAME) left at home without you for more than 1 hour? _____ If the answer is 0 times >> skip to question 542
541	¿Who does usually stay in charge of (NAME) in these cases?	1	Nobody / he/she stays alone
		2	He stays with > indicate code _____ 2. (NAME)'s mother; 3. (NAME)'s father; 4. (NAME)'s brothers or sisters; 5. (NAME)'s grandparents; 6. Other
542	Sometimes children misbehave and parents have different ways of managing discipline. During the last 7 days, when (NAME) misbehaved, how did you disciplined him/her? Enumerator > do not read the options, mark all that apply	1	Slapped, pinched child or used "chicote"
		2	Shouted at child
		3	Threatened child
		4	Child was left alone until he/she stopped crying
		5	Talked to child
		6	Other (specify) _____
		7	Did nothing
543	During the last 7 days, when (NAME) was well-behaved or did something you asked, what did you do? Enumerator > do not read the options, mark all that apply	1	Gave him/her a reward
		2	Hugged, cuddled him/her
		3	Praised him/her verbally, applauded him/her
		4	Did nothing
		5	Other (specify) _____
B Expectations regarding child development			
> Let us now talk about the future of (NAME)			
550	What job would you like (NAME) to do in the future? Which educational attainment would you like (NAME) to achieve?		Specify: _____ Primary ☐ 1 Second. ☐ 2 Higher ☐ 3
551	Let us now think about the path that (NAME) needs to follow in order to achieve these goals. How much do you agree or disagree with the following statements?	1	Now that (NAME) is a small child is an important stage if he/she is to achieve (...) education and become (...)
		2	It does not depend on you that (NAME) achieves (...) education and becomes (...)
		3	(NAME) will start learning only after he/she enrolls in school
		4	If you spend more time playing with (NAME), you will help him/her to achieve (...) education and become (...)
		5	You are more important than school teachers for (NAME) to be able to achieve (...) education and become (...)
		6	Now that (NAME) is a small child it is better if he/she learns on his/her own
		7	It is useless to tell stories to (NAME) because he/she does not understand them
		8	Now that (NAME) is a small child he only needs to be clean and well fed

C Interactions during the interview, physical appearance and hygiene					
560	Was it possible to observe the interaction between the caregiver and the child during the interview?	1	Yes		
		2	No >> state the reason _____ >> skip to question 562		
			Yes	No	
561	Enumerator: choose yes/no according to what you observed regarding the interaction of the caregiver and his/her partner with the child during the interview	1	Gave toys or appropriate activities to the child	1	2
		2	Child was kept under their supervision	1	2
		3	They spontaneously vocalized to child at least twice	1	2
		4	They responded verbally to child's verbalizations at least once	1	2
		5	They told the child the name of an object or person at least once	1	2
		6	They spontaneously praised child at least twice	1	2
		7	Their voice conveyed positive feelings towards child	1	2
		8	They caressed/kissed/hugged child at least once	1	2
		9	They did not shout to child	1	2
		10	They did not express annoyance or hostility to child	1	2
		11	They did not slap or spank child	1	2
		12	They did not scold or criticize child	1	2
		13	They did not interfere or restrict child more than three times	1	2
562	Enumerator: choose yes/no according to what you observed regarding the physical appearance and hygiene of the caregiver during the interview	1	His/her face is extremely dirty	1	2
		2	His/her hair is extremely dirty	1	2
		3	His/her hands are dirty	1	2
		4	His/her clothes are extremely dirty	1	2
		5	Has open wounds	1	2
563	Was it possible to observe the physical appearance and hygiene of the child during the interview?	1	Yes		
		2	No >> state reason _____ >> finish questionnaire		
			Yes	No	
564	Enumerator: choose yes/no according to what you observed regarding the physical appearance and hygiene of the child during the interview	1	His/her face is extremely dirty	1	2
		2	His/her hair is extremely dirty	1	2
		3	His/her hands are dirty	1	2
		4	His/her clothes are extremely dirty	1	2
		5	Has appropriate clothes for the weather	1	2
		6	Nappies (if any) have not been changed for a long time	1	2
		7	Has open wounds	1	2

Figure 4B
Characterization of the quality of the home environment



Appendix 4.5

Complete regression results for treatment effects on the quality of the home environment

	Aggregate index		Index components			
	All components (1)	Complete information (2)	Play activities (3)	Responsiveness (4)	Play material (5)	Care and hygiene (6)
Treatment	0.059*** (0.013)	0.057*** (0.014)	0.068*** (0.013)	0.062*** (0.021)	0.098*** (0.023)	0.010 (0.011)
Baseline index	0.244*** (0.048)	0.235*** (0.039)	0.195*** (0.035)	0.062 (0.043)	0.102** (0.047)	0.114*** (0.029)
Household wealth index (0-1)	0.140*** (0.039)	0.080* (0.042)	0.080 (0.049)	0.057 (0.047)	0.256*** (0.078)	0.140*** (0.035)
Household members	0.001 (0.002)	-0.000 (0.003)	-0.004 (0.004)	0.001 (0.003)	0.006 (0.004)	-0.000 (0.001)
Caregiver's age	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Caregiver's educ. attainment (years)	0.004*** (0.001)	0.003 (0.002)	0.009*** (0.002)	0.003 (0.002)	0.008*** (0.003)	0.002 (0.001)
Caregiver is married (yes = 1)	0.003 (0.013)	0.004 (0.013)	0.004 (0.015)	-0.008 (0.020)	0.030 (0.021)	0.020* (0.010)
Hours worked last week by caregiver	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
Number of children under caregiver's care	-0.040* (0.020)	-0.034* (0.019)	-0.048* (0.028)	-0.022 (0.028)	-0.096** (0.044)	-0.020 (0.021)
Caregiver first language is Spanish (yes = 1)	0.002 (0.013)	0.012 (0.014)	0.029* (0.015)	-0.017 (0.018)	0.001 (0.024)	0.023* (0.012)
Caregiver is the mother (yes = 1)	-0.038 (0.032)	-0.044** (0.021)	-0.054 (0.039)	0.029 (0.046)	-0.095** (0.045)	-0.043*** (0.012)
Child's age (months)	-0.001* (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.002 (0.002)	0.001** (0.001)
Child is male (yes = 1)	0.000 (0.008)	0.011 (0.009)	-0.004 (0.015)	0.018* (0.010)	-0.006 (0.014)	0.012 (0.008)
Constant	0.483*** (0.055)	0.482*** (0.046)	0.338*** (0.072)	0.559*** (0.083)	0.491*** (0.081)	0.796*** (0.059)
Observations	486	614	614	548	537	558
R-squared	0.294	0.186	0.196	0.073	0.204	0.171

Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Chapter 5

Constraints Influencing the Effects of Home-Visiting Interventions on Parental Behaviour

Abstract

Home-visiting interventions implemented in the developing world have the explicit objective of changing caregiver behaviour so they offer better stimulation opportunities to their children. Despite this, little analysis has been devoted to understanding the mechanisms behind caregivers' behavioural change. Instead, the literature places a strong emphasis on the transfer of parenting skills as a key determinant of the intervention's effectiveness, assuming that the lack of these skills is the major binding constraint to be confronted. This chapter explores the constraints that limit the effect of a home-visiting intervention on caregivers' engagement in educational play activities with their children. I use data on parenting practices and caregiver beliefs regarding the importance of parent-child interactions for child development, collected from the control and treatment groups of families randomized for the evaluation of the rural component of programme *Cuna Mas* in Peru. Results reveal that: (i) low levels of wealth can render the intervention ineffective in changing parenting practices and the effect exhibits a positive gradient for most of the support of the wealth distribution; (ii) there is no evidence suggesting this heterogeneity is driven by an unequal transfer of parenting skills to caregivers; and (iii) caregivers' beliefs exhibit a pattern of heterogeneity similar to that displayed by their behavioural change. This suggests there are constraints besides the lack of parenting skills that determine the effect of the intervention and which seem related to the information available to caregivers regarding the importance of parent-child interactions for child development. I propose a simple time-allocation model to formalize these notions and offer a way of thinking about the mechanisms behind caregivers' behavioural change. By giving a role to the expected return to the time spent on childcare and relating this expectation to caregivers' beliefs regarding the importance of parent-child interactions, this model is capable of producing results consistent with the available evidence.

5.1 Introduction

Results discussed in the previous chapter show that a home visiting ECD intervention working at scale can provoke a change in caregiver behaviour conducive to the provision of more stimulation opportunities for her child. Using data from the rural home-visiting

component of programme *Cuna Mas* in Peru, it was possible to identify an increase in caregiver engagement in educational play activities with her child in addition to those occurring during the home visits. This effect is positive, statistically significant and has a size of 0.3 standard deviations.

In the previous chapter, I also highlighted the literature's lack of attention to the mechanisms behind caregivers' behavioural change and how its strong emphasis on parenting skills implicitly assumes that caregivers' lack of knowledge about certain activities that can be performed with their children and of the ability to engage in them are the major binding constraint. I argued this assumption can be rather strong especially if we are working with poor families who, by definition, face scarcity in many dimensions.

In this chapter I propose an analysis of the potential constraints limiting the effect of the abovementioned intervention on caregiver participation in play activities with their children. The objective is to determine if the available evidence is consistent with poor parenting skills being the main binding constraint limiting caregiver participation in play activities or if the evidence suggests there are other limiting factors.

I start from the premise that treatment is transferring skills to caregivers that enable them to engage in more educational play activities with their children. Therefore, if caregivers exhibit a heterogeneous incremental engagement due to the intervention, it is either because there has been an unequal transfer of skills or because they respond differently to these skills. If the lack of parenting skills is the major constraint preventing caregivers from engaging in more play activities, heterogeneity in treatment effects must respond to the first possibility. In other words, some caregivers have ended up with more additional skills than others and this has determined different treatment effects in terms of their engagement in play activities. If there are other elements besides skill conditioning caregivers' response,

then heterogeneity in treatment effects can be present even if the shift in parenting skills has been homogeneous across caregivers. In this case, a similar pattern of heterogeneity should be present in those elements conditioning caregivers' response.

Based on these notions, the analysis is organised in three parts. Section 5.2 presents empirical evidence about the presence of heterogeneous treatment effects by caregivers' wealth and educational attainment. Household wealth provides a broad view of the degree of scarcity faced by caregivers and their families. A positive wealth gradient in treatment effects would suggest that the intervention is requiring parents to input resources of their own that poorer caregivers are finding more difficult to forgo.

A positive wealth gradient in treatment effects can also be the consequence of an unequal shift in parenting skills if the process of skill acquisition is influenced by some variable that correlates with wealth. Caregivers' educational attainment usually correlates with wealth and can determine the effectiveness with which new information and knowledge regarding educational play activities is transmitted to them. Therefore, the first part of the analysis also evaluates whether treatment effects differ by caregivers' education in an effort to produce evidence consistent with an unequal transfer of parenting skills to caregivers.

The second part of the analysis is presented in Section 5.3 and discusses results obtained when measuring caregivers' beliefs about the importance of parent-child interactions for child development, and presents treatment effects on this variable. This is one of the novel features of the analysis which aims at exploring if these beliefs play a role as a conditioning factor for caregivers' response to the transfer of parenting skills. As already discussed, conditioning factors should exhibit a similar pattern of heterogeneity to treatment effects on caregiver behaviour. The third part of the analysis (in Section 5.4) presents a simple model to rationalise the empirical findings and provide a way of thinking about the mechanisms

behind caregivers' behavioural change. Finally, Section 5.5 presents some concluding remarks.

5.2 Heterogeneity of treatment effects: the role of caregivers' wealth and educational attainment

Three different techniques are employed to determine whether the effects of the intervention differ depending on caregivers' wealth and educational attainment. The first consists in estimating different treatment effects for different groups within the wealth and education distributions. If we divide these distributions into K groups, this can be easily accomplished using the following specification:

$$y_i = \alpha_1 + \beta_1 T_i + \beta_2 T_i G_{i2} + \dots + \beta_k T_i G_{ik} + \alpha_2 G_{i2} + \dots + \alpha_k G_{ik} + x_i' \gamma + \varepsilon_i \quad (5.1)$$

Where T_i denotes individual i treatment status, G_{ik} indicates membership to group k , and x_i' is a vector of pre-treatment controls. The treatment effect for the first group is given by $\hat{\beta}_1$, while the treatment effect for group k is given by $\hat{\beta}_1 + \hat{\beta}_k$. This specification has the advantage of allowing one to directly assess the significance of the difference of treatment effects with respect to the first group by simply looking at the significance of coefficients $\hat{\beta}_2, \dots, \hat{\beta}_k$.

The second technique allows wealth and education to affect treatment effects in a continuous way. This is tested by introducing an interaction between treatment status and caregivers' wealth or educational attainment. Formally:

$$y_i = \alpha_1 + \lambda_1 T_i + \lambda_2 T_i z_i + x_i' \gamma + \varepsilon_i \quad (5.2)$$

$$y_i = \alpha_1 + \lambda_1 T_i + \lambda_2 T_i z_i + \lambda_3 T_i z_i^2 + x_i' \gamma + \varepsilon_i \quad (5.3)$$

Where z_i refers to caregivers' household wealth or educational attainment, depending on the source of heterogeneity under analysis. In specifications (5.2) and (5.3), the estimated effect of the treatment is given by $\hat{\lambda}_1 + \hat{\lambda}_2 z_i$ and $\hat{\lambda}_1 + \hat{\lambda}_2 z_i + \hat{\lambda}_3 z_i^2$, respectively. Specification (5.3) allows for a non-linear relationship between z_i and the treatment's impact.

Finally, a non-parametric fit of the relationship between the variety of play activities offered by the caregiver and her educational attainment or household wealth provides an even more flexible way of evaluating the presence of heterogeneity in treatment effects. For this, we need to assess the difference between the non-parametric fit built using data from the control and treatment groups. Systematic changes in this difference will be indicative of heterogeneity in treatment effects. If we suspect treatment effects exhibit a positive or negative gradient according to a certain variable, the behaviour of this difference can serve to evaluate the robustness of such gradient.

Tables 5.1 and 5.2 and Figure 5.1 present the results obtained from the three techniques explained above. Estimations following specification (5.1) were obtained after dividing the wealth and education distributions into four groups (quartiles). Results are presented in Table 5.1. Table 5.2 reports estimates from specifications (5.2) and (5.3). Panel A in Figure 5.1 shows the behaviour of treatment effects according to caregiver wealth and education, estimated from the non-linear specification given in (5.3).⁹⁴ Finally, Panel B in Figure 5.1 presents a non-parametric fit by treatment status of the relationship between play activities and caregiver wealth and educational attainment.

⁹⁴ Upper and lower bounds for the x-axis variables were set at the mean values in the first and fourth quartiles of the wealth and education distributions.

Table 5.1

Treatment effects by household wealth and caregiver education quartiles on the variety of play activities offered by the caregiver at least twice per week

(a) Household wealth				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.207 (0.186)	0.461* (0.268)	1.131*** (0.240)	0.239 (0.239)
Difference w.r.t Q1		0.254 (0.316)	0.924*** (0.288)	0.031 (0.305)
(b) Caregiver education				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.439** (0.174)	0.621** (0.272)	0.591* (0.335)	0.632** (0.241)
Difference w.r.t Q1		0.182 (0.326)	0.152 (0.386)	0.193 (0.273)

Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions include pre-treatment controls.

Table 5.2

Interactions between treatment status, household wealth and caregivers' educational attainment


	z(i) = household wealth		z(i) = caregiver's educational attainment	
	<i>Linear</i>	<i>Non-linear</i>	<i>Linear</i>	<i>Non-linear</i>
Treatment	0.251 (0.329)	-1.210 (0.843)	0.411* (0.213)	0.331 (0.252)
Treatment*z(i)	0.611 (0.691)	7.650* (3.956)	0.019 (0.031)	0.0765 (0.110)
Treatment*z(i)^2	--	-7.553* (4.115)	--	-0.005 (0.009)
z(i)	0.781* (0.406)	-2.320 (2.184)	0.097*** (0.022)	-0.069 (0.059)
z(i)^2	--	3.193 (2.028)	--	0.013*** (0.005)

Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

All regressions include pre-treatment controls.

Figure 5.1
Treatment effects by household wealth and caregiver educational attainment



Panel A depicts estimated treatment effects ($\hat{\lambda}_1 + \hat{\lambda}_2 z_i + \hat{\lambda}_3 z_i^2$) and 95% confidence intervals obtained from the following specification: $y_i = \alpha_1 + \lambda_1 T_i + \lambda_2 T_i z_i + \lambda_3 T_i z_i^2 + x_i' \gamma + \varepsilon_i$, where z_i refers to caregivers' household wealth (left hand side graph) and educational attainment (right hand side graph). Panel B depicts a locally weighted scatterplot smoothing of play activities on household wealth and caregiver education. In each graph, results are presented for the control and treatment groups. Therefore, treatment effects are given by the vertical distance between the two functions.

Several important results emerge from the analysis presented in the tables and figures above.

- (i) Household wealth has a significant effect on the impact of the intervention.
- (ii) There is evidence suggesting that treatment effects are a non-linear function of wealth. There is, however, a monotonic and significant increase in treatment

effects as household wealth raises, up to the 75th percentile of the wealth distribution.

- (iii) There is no evidence of significant treatment effects in the lower 25% of the wealth distribution. This result is robust to all the specifications considered.
- (iv) There is no evidence of heterogeneity in treatment effects on the variety of play activities by caregivers' educational attainment. In particular, impacts are positive, significant and amount to approximately 0.5 activities across the entire education distribution.

5.3 Measurement and treatment effects on caregivers' beliefs regarding the importance of parent-child interactions for child development

Caregivers' beliefs regarding the importance of parent-child interactions for the development of their children play an important role in shaping parenting practices. Under the assumption that parents try to provide what they believe is best for their children, variation in these beliefs should conduce to variation in observed behaviour. As discussed in Bornstein and Putnick (2012), "variation on childrearing philosophies, values and beliefs mediates differences in childrearing practices vis-à-vis local and larger physical and social environments" (Bornstein and Putnick (2012), p. 57). Caregivers' understanding of the significance of their role and of particular activities that can be carried out during their children's early years are part of these set of beliefs (Hoff et al., 2002).

Cuna Mas has no explicit objective in terms of caregivers' beliefs about the importance of parent-child interactions for child development. Despite this, *facilitadoras* are expected to offer messages about the importance of parenting practices for the child's wellbeing (PNCM, 2013). It is, therefore, reasonable to postulate that the intervention can affect caregivers' beliefs regarding their role and the importance of certain activities for child

development. In fact, this effect could be one of the channels through which behavioural change is achieved. More in general, these beliefs can condition caregivers' response to the transfer of parenting skills.

As described in the previous chapter, this study aimed at collecting information on caregivers' beliefs about the importance of parent-child interactions for child development by means of an eight-item scale. Each item proposed a statement and the caregiver had to indicate (through a visual aid) to what extent she agrees or disagrees with that statement. In order to ensure that caregivers share a meaningful and concrete idea of "child development" when completing the scale, they were first asked about the ideal educational attainment and occupation they would like their child to achieve.⁹⁵ The statements of the scale make reference to these goals and how their accomplishment relates to early childhood as a sensitive period for development, caregivers' role during this period, and caregivers' engagement in educational play activities. The specific items are presented in Appendix 4.4.

Results in terms of the internal consistency and validity of this instrument are mixed. On one hand, the scale exhibited low internal consistency with a Cronbach's alpha estimate that ranged between 0.46 (with the full sample) and 0.5 (if we focus on caregivers with more than completed primary education -45%-, or those who have Spanish as their mother tongue -64%-).

During the pilot exercise, it was noted that caregivers had difficulties understanding the statements and the nature of the exercise (i.e. that they had to communicate an opinion regarding what was being said). These difficulties can be related to caregivers' limited

⁹⁵ The concept of "child development" can be rather abstract and different caregivers might have different views about it. Because of this, the ideal educational attainment and occupation that caregivers' wish for the child were proposed as a concrete and common outcome when answering the scale.

verbal communication skills, as confirmed by the fact that the internal consistency of the instrument is sensitive to the educational attainment of the participant. Caregivers' mother tongue also played a role, possibly because of noise introduced when the statements were translated to the local idiom (Quechua).

After the pilot exercise, efforts were made to simplify the statements and train enumerators to guide the process. As a result, seven out of the original eight items exhibited the expected correlation with the rest⁹⁶ but, as already mentioned, overall consistency ended up being low (Cronbach alpha = 0.5), although still above the range of values judged as “unacceptable” (Gliem and Gliem, 2003).

On the other hand, in terms of validity, the scores obtained from the final 7-item scale exhibit the expected positive socioeconomic gradient and, importantly, also predict the scores obtained in the quality of the home environment index. Table 5A in Appendix 5.1 documents the relationship between the scores that can be obtained from the scale,⁹⁷ household wealth and caregiver education. We encounter a positive relationship⁹⁸ consistent with what has already been documented for the beliefs parents hold about their role in achieving the goals they value for their children: lower-SES parents believe they have less control over the outcome and should play a more passive role (Bornstein and Putnick, 2012; Hoff et al., 2002).

Tables 5B, 5C and 5D in Appendix 5.1 present results for the partial correlation between the scores obtained in the scale and those obtained in the quality of the home environment

⁹⁶ The last item was excluded because it exhibited a negative correlation with the rest.

⁹⁷ The two outcome measures are: (i) an average score (with values between 1 and 4); and (ii) an aggregate score (with values between 7 and 28). The aggregate score is built using the information of those caregivers that answered all seven items.

⁹⁸ One could suspect that the positive socioeconomic gradient in beliefs is driven by a positive socioeconomic gradient in aspirations. Although, in principle, they reflect different concepts, the questionnaire could have failed to capture them separately. There is, however, no evidence of this as 99% of caregivers share the same aspiration in terms higher education attainment for their children (i.e. there is no socioeconomic gradient in aspirations).

index, its subcomponents and the variety of play activities. There is a positive and statistically significant partial correlation (after holding observable pre-treatment characteristics constant) between our measure of caregivers' belief about the importance of parent-child interactions for child development and caregivers' parenting practices. Remarkably, this positive partial correlation is significant for the two subcomponents more closely related to caregivers' actions and attitudes towards their children ("interaction and play activities" and "responsiveness and control practices"). Within the first subcomponent, caregiver beliefs are also a good predictor of the variety of play activities offered.

In Table 5.3 below, I present treatment effects for the two outcome measures built with the final 7-item scale. There is no evidence of significant average treatment effects on caregivers' beliefs regarding the importance of parenting activities for child development. An analysis by wealth and education groups, however, reveals the existence of positive and significant effects among the wealthiest caregivers (see Table 5.4). A comparison between treatment and control groups of the relationship between caregivers' beliefs and household wealth (see Figure 5.2) confirms the presence of consistently larger treatment effects in the upper part of the wealth distribution. This same pattern is not apparent if we evaluate treatment effects by caregivers' education. These results are robust to considering the aggregate score of the scale (see Table 5E in Appendix 5.1) and to restricting the sample to those caregivers that have Spanish as their mother tongue, for whom the internal consistency of the scale is larger (see tables 5F, 5G and 5H in Appendix 5.1).

Table 5.3

Treatment effects on caregivers' beliefs regarding the importance of parent-child interactions for child development

	Obs.	Mean control group	Mean treatment group	Difference	Controlled difference	d
Average score (1-4)	614 (C=320; T=294)	2.951	2.995	0.044 (0.039)	0.034 (0.032)	0.097
Aggregate score (7-28)	579 (C=304; T=275)	20.776	21.076	0.300 (0.266)	0.219 (0.230)	0.089

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Controlled difference is the estimated effect including pre-treatment controls.

Effect size (d) calculated as the standardized controlled difference.

Table 5.4

Treatment effects by household wealth and caregiver education on caregivers' beliefs regarding the importance of parent-child interactions for child development (average score)

(a) Household wealth				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	-0.013 (0.047)	0.058 (0.051)	-0.056 (0.060)	0.144*** (0.058)
Difference w.r.t Q1		0.070 (0.060)	-0.043 (0.085)	0.157** (0.075)
(b) Caregiver education				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.025 (0.049)	0.126*** (0.044)	0.057 (0.057)	-0.045 (0.077)
Difference w.r.t Q1		0.101 (0.068)	0.032 (0.076)	-0.070 (0.088)


Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions include pre-treatment controls.

Figure 5.2

Relationship between caregivers' beliefs regarding the importance of parent-child interactions for child development (average score) and caregivers' household wealth and education



Figures depict a locally weighted scatterplot smoothing of the average score obtained in the beliefs scale on household wealth (panel A) and caregiver education (panel B). In each graph, results are presented for the control and treatment groups. Therefore, treatment effects are given by the vertical distance between the two functions.

5.4 A simple time allocation model to explore the mechanisms

In the previous sections I have shown that the intervention has caused a heterogeneous shift in caregiver behaviour. In fact, the size of treatment is positively related to caregivers' wealth and is not significant among the poorest caregivers. In addition, I have not found evidence that this heterogeneity is driven by an unequal shift in parenting skills. Caregivers' educational attainment correlates with wealth and can condition the transmission of skills, but treatment effects are homogenous across the entire caregiver education distribution.⁹⁹ Finally, caregivers' beliefs regarding the importance of parent-child interactions for child

⁹⁹ An unequal acquisition of parenting skills by caregivers could also be caused by different intensities of treatment. In this regard, there is no evidence of heterogeneity in the intensity of treatment by caregiver wealth. In particular, the number of home-visits received until the follow-up survey is equal to 25 across the entire wealth distribution. The expansion in play materials is also homogenous across the wealth distribution (see Appendix 5.2). Also recall that treatment effects on the variety of play activities are insignificant among the poorest caregivers. Given that these caregivers have received the same materials and visits than the rest, it is implausible that the absence of effect is entirely explained by a failure to convey parenting skills.

development exhibit a positive wealth gradient as well as positive treatment effects concentrated in the upper 25% of the wealth distribution.

In this section I aim at rationalising these findings by means of a simple time allocation model. With this model, I also seek to provide a way of thinking about the mechanisms behind the effects of parent-child focused interventions on caregivers' engagement in cognitively stimulating activities with their children.

The model rests on the following premises: (i) treatment transfers parenting skills to caregivers and can also shift their beliefs regarding the importance of parent-child interactions; (ii) parenting skills allow caregivers to engage in more educational play activities with their children; (iii) from the point of view of caregivers, parenting skills are given, which means that these skills cannot be purchased in the market;¹⁰⁰ and (iv) to produce play activities caregivers also need to input time.

Regarding this last assumption, the literature offers evidence showing that family wealth, parental time devoted to childcare and parenting activities relevant for child development are all positively related. In fact, it has been documented that the amount of time that parents devote to their children in developed and developing countries increases with families' socioeconomic status (Guryan et al., 2008) and there is strong causal evidence linking the amount of time that parents and children spend together in play activities to child development (Fiorini and Keane, 2014). Based on this evidence and the nature of the outcome variable under analysis (the number of types of play activities offered to the child

¹⁰⁰ Absence of a market for these skills is a reasonable assumption given that the supply of play groups or nurseries is very limited in rural areas. In fact, it is this limited supply what *Cuna Mas* seeks to tackle.

at least twice per week) it is reasonable to assume that parenting skills need to be complemented with caregiver time in order to produce a shift in caregiver behaviour.¹⁰¹

One could argue that the effect of parenting programmes comes from making more productive (in terms of child development) the time that parents already spend with their children. An important implication of this assumption is that parents would always be willing to take up programme activities as these provide a time allowance: caregivers are able to produce the same amount of child development with less time devoted to childcare. In this case, the only input needed to provoke a change in caregiver behaviour is a successful transference of parenting skills. This is the assumption implicitly made by the literature surveyed in the previous chapter and which I argue is quite strong given that poorer families, which appear less responsive to these interventions, face deprivations in many dimensions.

In the analysis that follows I prefer to assume that parents need to input time in order to produce more parenting activities. As already mentioned, this assumption is based on the nature of the outcome variable under analysis and the evidence presented in the literature. The outcome variable is the number of types play activities offered to the child for a given frequency and does not refer to the quality of activities that are already carried out by parents. Notice that a new type of play activity (e.g. sharing a book) could be produced without devoting more time to childcare if caregivers were already engaged in activities with their children but these do not fall within those considered relevant for child development. This assumption is questionable. If the lack of early stimulation in low-SES

¹⁰¹ More types of educational play activities for a given frequency could be produced without devoting more time to childcare if we assume caregivers are already engaged in activities with their children but these are irrelevant for child development. This assumption is questionable. If the lack of early stimulation in low-SES families were due to the fact that parents are devoting time to their children but offering activities that are not appropriate for child development, we would not observe the positive SES gradient in childcare time nor the positive relation between the amount of time parents devote to children and child outcomes documented in the literature.

families were due to the fact that parents are devoting time to their children but offering activities that are not appropriate for child development, we would not observe the positive SES gradient in childcare time nor the positive relation between the amount of time parents devote to children and child outcomes documented in the literature.

The model proposed treats parenting skills as an input that enables a certain behaviour (in this case, caregiver engagement in educational play activities). Behavioural change, however, is not guaranteed because in addition to being able to engage in this behaviour, caregivers also have to understand it as something desirable. Under the logic of an economic model of rational behaviour, the latter will depend on the expected costs and benefits of allocating more time to childcare activities. Importantly, I will assume that the expected benefits of such decisions are reflected by caregivers' beliefs regarding the importance of parent-child interactions and are, therefore, susceptible to being affected by the intervention.

Another important implication of postulating that parenting skills and time are complementary to each other and that there is no market for parenting skills is that the actual amount of time allocated to childcare can be below the amount that equalizes marginal returns across all competing activities. As will be explained below, this situation is consistent with the existence of a binding constraint in parenting skills.

The model follows Kimmel and Connelly (2007) and Guryan et al. (2008) time allocation frameworks, but extends their original formulation to describe more explicitly the process through which childcare time is transformed into child-services. This is done by introducing a "parenting technology". As in Cunha et al. (2013), I also introduce uncertainty in caregivers' choices. In particular, caregivers do not know for certain the effect of their parenting activities on child outcomes. As already explained, the expected benefits of these actions are reflected by their beliefs regarding the importance of parent-child interactions.

Consider a static model where caregivers value consumption (C) and child-services (S). In principle, the concept of child services is broad enough to encompass the possibility that parents derive process utility from interacting with their children, and also outcome utility related to their children's human capital. For this analysis I will focus on the latter, which is true if parents are altruistic and care about the children's future consumption, and also if they are selfish and care about future transfers that can be received from their children.

Consumption goods can be produced by combining time devoted to home production (t_{HP}) and market purchases (X). To ease manipulation, assume that child's human capital depends exclusively on parenting practices (PA). These practices comprise the educational play activities analysed in previous sections and are produced by combining time devoted to childcare (t_C) and parenting skills (PS). The caregiver can also devote time to market production (t_{MP}) in exchange for a wage (w).

Because caregivers face uncertainty regarding the outcome of their parenting practices, they maximize expected utility conditional on the information available to them at the time choices are made (I). The problem can be stated as follows.

$$\max E(U|I) = U(C, E(S|I)) \quad (5.4)$$

Given:

$$C = C(t_{HP}, X; \theta) \quad (5.5)$$

$$E(S|I) = S(PA; E(\phi|I)) \quad (5.6)$$

$$PA = P(t_C, PS) \quad (5.7)$$

And subject to:

$$T = t_{HP} + t_{MP} + t_C \quad (5.8)$$

$$P_X X = wt_{MP} + V \quad (5.9)$$

Equation (5.5) indicates the technology for the production of consumption goods, where θ is an efficiency parameter. Equations (5.6) and (5.7) play an important role in this analysis and describe the processes and technologies involved in the production of child human capital. Function $S(\cdot)$ in equation (5.6) can be viewed as a simplified version of an early childhood skill formation technology.¹⁰² Parameter ϕ denotes the marginal return to parenting activities in terms of child's human capital. Caregivers' uncertainty regarding the returns to parental investments in children is captured through the expected value of this parameter for a given set of information $(E(\phi|I))$.

The function indicated in (5.7) captures the parenting technology: time and parenting skills are combined to produce an activity (e.g. sharing a book with the child). Consistent with the definition used in the preceding sections, parenting skills refer to caregivers' knowledge of particular educational play activities that can be carried out with their children and the ability to perform them. Finally, equations (5.8) and (5.9) define the time and monetary budget constraints, respectively. In (5.9), variable V captures other sources of income.

In this model, the allocation of time between the three competing activities (childcare, market production and home production) will depend on preferences regarding consumption goods and child human capital, and caregivers' productivity in the goods and child services sectors. Under this setting, a home-visiting intervention can be viewed as one that seeks to increase the productivity of caregivers in the child services sector in order to promote better child outcomes. This is pursued by increasing parenting skills, which means shifting PS .

¹⁰² Since this analysis focuses on early childhood (ages 0-3) it is reasonable to assume that the main inputs for skill formation during this transition are provided at home.

One key feature governing the outcomes of the model is the degree of complementarity assumed between caregiver time and skill for the production of PA . As explained above, this analysis will assume there is high degree of complementarity. An important implication of this assumption is that we understand the home-visiting intervention as one that provides skills to caregivers that allow them to engage in time-consuming activities that improve child outcomes.

Consistent with the abovementioned assumption, let the parenting technology be given by: $PA = \min(\varphi t_C, \gamma PS)$ where, as already explained, the amount of parenting skills applied is not a choice variable. In addition, assume that all the other functions are concave and continuously differentiable. Because PS is not a choice variable, the observed amount of childcare time will depend on whether parenting skills represent a binding constraint or not.

Let us start considering the case where PS are non-binding. Assume that PS are sufficiently large such that, for a given choice of t_C , $\min(\varphi t_C, \gamma PS) = \varphi t_C$. Under these circumstances, the optimal amount of time devoted to the production of PA (t_C^*) will be such that the expected marginal returns are equated across the three time use alternatives. This implies that (t_C^*) must satisfy the following condition:

$$\begin{aligned} U_S^* [S(P(t_C^*, PS); E(\phi|I)) - S(P(t_C^*, PS); E(\phi|I))] \\ = U_C^* \left[C \left(t_{HP}^*, \left(\frac{w}{P_X} \right) t_{MP}^* + \frac{V}{P_X}; \theta \right) - C \left(t_{HP}^{*'}, \left(\frac{w}{P_X} \right) t_{MP}^{*'} + \frac{V}{P_X}; \theta \right) \right] \end{aligned} \quad (5.10)$$

where U_S^* and U_C^* are the marginal utilities of child human capital and consumption, respectively, evaluated at the time allocations ($t_C^*, t_{HP}^*, t_{MP}^*$). This expression states that the optimal amount of time devoted to childcare will be the amount that guarantees that the expected utility gain from shifting time towards childcare to produce one more activity (at the left hand side of the equation, where $t_C^{*'} - t_C^* = 1/\varphi$) is equal to the utility loss

produced by the fall in consumption resulting from the decline in time devoted to either home or market production (at the right hand side of the equation, where $t_{HP}^{*'} \leq t_{HP}^*$, $t_{MP}^{*'} = T - t_C^{*'} - t_{HP}^{*'}$).

A situation where PS are binding means that for a given PS and t_C^* , $\min(\phi t_C^*, \gamma PS) = \gamma PS$. Notice that, in this case, the maximum amount of time that could be rationally devoted to producing PA (t_C^{max}) will be such that $\phi t_C^{max} = \gamma PS$,¹⁰³ which means that:

$$t_C^{max} = \gamma PS / \phi \quad (5.11)$$

If we combine both scenarios, we obtain the following solution for the actual time devoted to childcare:

$$t_C = \begin{cases} t_C^* & \text{if } \min(\phi t_C^*, \gamma PS) = \phi t_C^* \\ t_C^{max} & \text{if } \min(\phi t_C^*, \gamma PS) = \gamma PS \end{cases} \quad (5.12)$$

Let us now use the solution to this model to explore the effects of the home-visiting intervention on the production of PA and, in particular, to offer an explanation for the empirical evidence presented in the previous sections. For this, denote pre-treatment parenting skills as PS^C and assume that treatment effectively shifts these skills to a level $PS^T > PS^C$.

Two key elements determining the effect of this shift on the production of PA are the pre-treatment difference between t_C^* and t_C^{max} , and the possibility of a change in some other determinant of caregiver behaviour being induced by the intervention. As already mentioned, I will consider the possibility of a shift in the expected return to parenting activities ($E(\phi|I)$).

¹⁰³ If PS are binding then $PA = \min(\phi t_C^*, \gamma PS) = \gamma PS$. Allocating any amount of t_C such that $\phi t_C > \gamma PS$ would not be consistent with the utility maximizing principle of the model because part of the time devoted to PA is not contributing to its production but has an opportunity cost.

In what follows, I will consider two scenarios in terms of the abovementioned determinants of treatment effects, and contrast how these scenarios fare in explaining the evidence. The first scenario will illustrate a situation where caregivers' lack of parenting skills is the major constraint preventing them from offering a more cognitively stimulating environment to their children. This is consistent with the emphasis given in the literature to the transfer of parenting skills when addressing the effectiveness of this type of interventions.

The second scenario will also give a role to parenting skills but will draw attention to the fact that these skills need to be complemented with time and that reallocating time towards childcare faces an opportunity cost and has an expected return. As will be discussed below, giving an explicit role to the expected return to childcare time in caregiver time allocation decisions will allow one to explain the available evidence on treatment effects in a more coherent way.

Panels A and B in Figure 5.3 illustrate these two scenarios. Both panels illustrate the relationship between childcare time and household wealth before and after treatment. Superscript C denotes childcare time before treatment and superscript T indicates post-treatment childcare time. The positive slopes are consistent with the positive SES gradient in childcare time already documented in the literature (Guryan et al., 2008).¹⁰⁴ Given the parenting technology and the solution for the amount of time devoted to childcare expressed in (5.12), there is a direct relationship between childcare time and the variety of play activities offered to the child by the caregiver (PA). In this regard, the positive wealth gradient in parenting activities implicit in Figure 5.3 is consistent with the evidence

¹⁰⁴ It is not the objective of this analysis to explain the positive wealth gradient in childcare time. In the model, the positive wealth gradient in t_C^{max} can be accommodated if we allow a positive correlation between household wealth and parenting skills (recall $t_C^{max} = \gamma PS/\varphi$). The positive wealth gradient in t_C^* can be accommodated if we allow greater household wealth to be accompanied by a sufficiently large expansion in V . Notice that, other things equal, an increase in V exogenously raises consumption levels, reduces the marginal utility of consumption and, thus, reduces the utility loss derived from allocating time away from market or home production and towards childcare.

presented in Figure 5.1 and has also been extensively documented in the literature (Hoff et al., 2002).

Panel A in Figure 5.3 presents the scenario where the transfer of parenting skills plays the main role in determining treatment effects. Accordingly, this first scenario only considers that treatment produces a positive shift in parenting skills so that $PS^T > PS^C$. In addition, it postulates a pre-treatment situation where $t_C^C = t_C^{Cmax} < t_C^{C*}$, so that a positive shift in parenting skills will suffice to induce a change in behaviour. Notice that as long as parenting skills are binding ($t_C^{C*} > t_C^{Cmax}$), caregivers will find it optimal to devote more time to childcare provided they are able to produce more activities. Therefore, panel A also introduces a sufficiently large difference between t_C^{C*} and t_C^{Cmax} so that parenting skills remain binding after treatment. In particular $t_C^T = t_C^{Tmax} < t_C^{C*}$.

Under this setting, the possibility of a wealth gradient in treatment effects depends entirely on there being a wealth gradient in the acquisition of parenting skills. In fact, the effect of treatment is given by $t_C^T - t_C^C = t_C^{Tmax} - t_C^{Cmax} = \gamma / \phi(PS^T - PS^C)$. As already discussed, however, there is no evidence of a heterogeneous transfer of skills that correlates with wealth. As shown in panel A of Figure 5.3, a homogenous increase in PS (depicted as a parallel shift in the t_C^{max} function) will not be able to produce a positive wealth gradient in treatment effects. To see this, notice that the effect of the intervention on childcare time in Figure 5.3 is given by the vertical distance between the two functions that define the area shaded in grey. Also recall there is a direct relationship between childcare time and PA .

Panel B in Figure 5.3 presents an alternative scenario where the increase in PS is accompanied by a positive shift in the expected return to childcare time caused by a change in the information set available to caregivers. In particular: $E(\phi|I^T) > E(\phi|I^C)$. This interpretation of the mechanisms behind the observed treatment effects also gives a role to


parenting skills and that is why skills are binding before treatment. The objective, however, is to draw attention to the expected return to childcare time and, therefore, the pre-treatment situation is represented as $t_C^C = t_C^{Cmax} = t_C^{C*}$. In other words, to produce more parenting activities caregivers need more skills but just having the skills will not imply that the outcome obtained from the additional childcare time is worth the forgone consumption.

Consistent with the evidence discussed above, the increase in PS is homogeneous across the wealth distribution (represented as a parallel shift in the t_C^{max} function) but the change in $E(\phi|I)$ is concentrated among the richest caregivers. This is represented by an increase in the slope of the t_C^* function.¹⁰⁵ As shown in panel B of Figure 5.3, this scenario is capable of producing a positive wealth gradient in treatment effects on the variety of play activities offered by the caregiver. This explanation postulates that the observed heterogeneity in caregivers' behavioural change is consistent with a homogenous shift in parenting skills because a similar pattern of heterogeneity is present in caregivers' beliefs regarding the importance of parent-child interactions, and these beliefs condition caregivers' response to the transfer of parenting skills.¹⁰⁶

¹⁰⁵ Notice that larger $E(\phi|I)$ directly increases the expected utility gain from shifting time towards childcare in (5.10) and therefore causes an increase in t_C^* .

¹⁰⁶ A sufficiently large slope in the t_C^* function prior to treatment $\left(\frac{\partial t_C^*}{\partial wealth} > \frac{\partial t_C^{Cmax}}{\partial wealth}\right)$ would also produce a positive wealth gradient in treatment effects without requiring a shift in caregiver beliefs. In this case, one could think of caregiver beliefs regarding the importance of parent-child interactions as a moderator of treatment effects. In Appendix 5.3, I present treatment effects differentiated according to the scores obtained in the caregiver beliefs scale. The evidence is consistent with beliefs conditioning caregivers' response to the transfer of parenting skills. In particular, bigger scores are accompanied by larger treatment effects on caregiver participation in educational play activities. Results, however, must be interpreted with caution as they consider post-treatment values of the beliefs scale. Correlation between these values and treatment status can introduce bias in the results.

Figure 5.3
Two scenarios for treatment effects on childcare time



5.5 Concluding remarks

Home-visiting ECD interventions implemented in the developing world have the explicit objective of enhancing caregivers' parenting practices so better stimulation opportunities are offered to their children. Despite this, little analysis has been devoted to understanding the mechanisms behind caregivers' behavioural change. Instead, the literature places a strong

emphasis on the transfer of parenting skills as a key determinant of the intervention's effectiveness. This requires the implicit assumption that caregivers' lack of knowledge about certain activities that can be performed with their children and of the ability to engage in them, are the major binding constraint to be confronted.

The analysis presented in this chapter aimed at contributing to the ECD literature by exploring the constraints that limit the effect of a home-visiting intervention on caregivers' behavioural change. In particular, the analysis focused on the home-visiting component of programme *Cuna Mas* currently working in rural Peru, and its effects on caregivers' engagement in educational play activities with their children and their beliefs regarding the importance of parent-child interactions for child development.

Results presented in this analysis indicate that there are elements besides a successful transfer of parenting skills that determine the effect of the intervention on caregiver behaviour and, thus, limit caregiver participation in play activities with their children. The evidence shows these constraints are related to household wealth and seem associated to the information available to caregivers regarding the importance of parent-child interactions for child development.

In this regard, I have found heterogeneous treatment effects on caregivers' engagement in play activities by household wealth. In fact, low levels of wealth can render the intervention ineffective in changing these behaviours, while the effect exhibits a positive gradient for most of the support of the wealth distribution. I have found no evidence suggesting this heterogeneity is driven by an unequal transfer of parenting skills to caregivers. In fact, caregivers' response is not conditioned by their educational attainment and the number of visits and the expansion in play materials are the same regardless of caregivers' wealth.

What differ according to caregivers' wealth are their beliefs regarding the importance of parent-child interactions for child development. The pattern of heterogeneity found in this data is similar to that exhibited by caregivers' behavioural change. In fact, caregivers' beliefs not only show a positive wealth gradient but also exhibit positive treatment effects among the richest caregivers. This evidence is consistent with caregivers' response to treatment being positively influenced by their beliefs regarding the importance of parent-child interactions.

A simple time-allocation model was proposed to formalize these notions and offer a way of thinking about the mechanisms behind caregivers' behavioural change. Importantly, the model proposes one to think about the parenting skills provided by the intervention as an enabling factor. These skills allow caregivers to produce more play activities provided they are willing to input time, which means that the effect of the intervention on caregiver behaviour depends on the expected costs and benefits of reallocating time to childcare.¹⁰⁷ By giving a role to the expected return to childcare time and relating this expectation to caregivers' beliefs regarding the importance of parent-child interactions, this model is capable of producing results coherent with the available evidence. The same is not true if one solely relies on parenting skills to explain caregivers' response to treatment.

¹⁰⁷ A different view is one in which the intervention is understood as providing skills that will increase the productivity of caregiver time for the production of play activities. The view proposed in this study is not only more coherent with the outcome variable under analysis (the number of types of play activities) but also consistent with recent evidence found for a similar intervention. In fact, a recent study by Attanasio et al. (2015) which analyses the effects of a home-visiting intervention in Colombia shows that child outcomes improved through an increase in the resources invested by parents in their children rather than through an increase in the productivity of these resources.

Appendix 5.1

Results related to caregivers' beliefs regarding the importance of parent-child interactions for child development

Table 5A

Socioeconomic gradients (by caregiver education and household wealth quartiles)

(a) Caregiver education	Q1	Q2	Q3	Q4
Caregiver beliefs regarding importance of parent child interactions (average score: 1-4)	2.835	2.95	3.004	3.164
Difference w.r.t Q1		0.115***	0.169***	0.329***
Caregiver beliefs regarding importance of parent child interactions (total score: 7-28)	20.021	20.715	21.016	22.164
Difference w.r.t Q1		0.694***	0.885***	2.143***

(b) Household wealth	Q1	Q2	Q3	Q4
Caregiver beliefs regarding importance of parent child interactions (average score: 1-4)	2.874	2.976	2.979	3.06
Difference w.r.t Q1		0.102**	0.105**	0.186***
Caregiver beliefs regarding importance of parent child interactions (total score: 7-28)	20.295	20.924	20.973	21.443
Difference w.r.t Q1		0.629**	0.678**	1.148***

. *** p<0.01, ** p<0.05, * p<0.1

Table 5B
Caregivers' beliefs partial correlation with the quality of the home environment (average score)

	Aggregate index		Index components			
	All components (1)	Complete information (2)	Play activities (3)	Responsiveness (4)	Play material (5)	Care and hygiene (6)
Treatment	0.057*** (0.013)	0.055*** (0.013)	0.065*** (0.013)	0.060*** (0.021)	0.097*** (0.023)	0.009 (0.011)
Baseline index	0.235*** (0.049)	0.222*** (0.040)	0.175*** (0.032)	0.055 (0.043)	0.103** (0.047)	0.112*** (0.030)
Household wealth index (0-1)	0.132*** (0.037)	0.074* (0.039)	0.071 (0.045)	0.046 (0.043)	0.253*** (0.077)	0.138*** (0.035)
Household members	0.001 (0.002)	-0.000 (0.002)	-0.005 (0.004)	0.000 (0.003)	0.006 (0.004)	-0.000 (0.001)
Caregiver's age	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Caregiver's educ. attainment (years)	0.003** (0.001)	0.002 (0.002)	0.006*** (0.002)	0.001 (0.002)	0.008** (0.003)	0.001 (0.001)
Caregiver is married (yes = 1)	-0.000 (0.013)	0.001 (0.013)	-0.002 (0.015)	-0.013 (0.020)	0.029 (0.021)	0.019* (0.010)
Hours worked last week by caregiver	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
Number of children under caregiver's care	-0.041** (0.019)	-0.034* (0.018)	-0.049 (0.029)	-0.022 (0.025)	-0.096** (0.043)	-0.019 (0.021)
Caregiver first language is Spanish (yes = 1)	-0.001 (0.013)	0.008 (0.014)	0.022 (0.015)	-0.022 (0.018)	-0.001 (0.024)	0.021* (0.012)
Caregiver is the mother (yes = 1)	-0.035 (0.032)	-0.043** (0.019)	-0.053 (0.036)	0.028 (0.045)	-0.094** (0.045)	-0.043*** (0.012)
Child's age (months)	-0.001 (0.001)	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.002 (0.002)	0.001** (0.001)
Child is male (yes = 1)	-0.000 (0.007)	0.010 (0.009)	-0.004 (0.014)	0.017* (0.010)	-0.006 (0.014)	0.012 (0.008)
Scale average score	0.043*** (0.012)	0.053*** (0.018)	0.089*** (0.018)	0.067*** (0.020)	0.021 (0.020)	0.016 (0.010)
Constant	0.377*** (0.066)	0.352*** (0.058)	0.115 (0.092)	0.397*** (0.099)	0.436*** (0.098)	0.756*** (0.060)
Observations	486	614	614	548	537	558
R-squared	0.309	0.205	0.222	0.094	0.206	0.174

Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Table 5C**Caregivers' beliefs partial correlation with the quality of the home environment (aggregate score)**

	Aggregate index		Index components			
	All components (1)	Complete information (2)	Play activities (3)	Responsiveness (4)	Play material (5)	Care and hygiene (6)
Treatment	0.057*** (0.013)	0.055*** (0.014)	0.067*** (0.014)	0.059*** (0.021)	0.098*** (0.023)	0.011 (0.010)
Baseline index	0.235*** (0.052)	0.224*** (0.042)	0.181*** (0.033)	0.035 (0.045)	0.104** (0.046)	0.115*** (0.034)
Household wealth index (0-1)	0.141*** (0.039)	0.080* (0.041)	0.080* (0.046)	0.056 (0.046)	0.261*** (0.081)	0.141*** (0.039)
Household members	0.002 (0.003)	0.001 (0.003)	-0.004 (0.004)	0.002 (0.003)	0.007 (0.004)	-0.000 (0.001)
Caregiver's age	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.002 (0.001)	-0.001 (0.001)
Caregiver's educ. attainment (years)	0.004*** (0.001)	0.002 (0.002)	0.006*** (0.002)	0.002 (0.002)	0.009*** (0.003)	0.001 (0.001)
Caregiver is married (yes = 1)	-0.000 (0.014)	0.002 (0.014)	-0.003 (0.016)	-0.011 (0.021)	0.030 (0.022)	0.019* (0.011)
Hours worked last week by caregiver	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
Number of children under caregiver's care	-0.039** (0.019)	-0.032* (0.018)	-0.047 (0.030)	-0.019 (0.025)	-0.093** (0.043)	-0.021 (0.021)
Caregiver first language is Spanish (yes = 1)	0.002 (0.013)	0.010 (0.015)	0.026 (0.016)	-0.018 (0.018)	0.007 (0.025)	0.019 (0.012)
Caregiver is the mother (yes = 1)	-0.036 (0.033)	-0.043** (0.019)	-0.054 (0.037)	0.029 (0.044)	-0.103** (0.048)	-0.042*** (0.012)
Child's age (months)	-0.001* (0.001)	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.003 (0.002)	0.001*** (0.001)
Child is male (yes = 1)	0.003 (0.009)	0.013 (0.010)	0.001 (0.016)	0.021* (0.011)	-0.002 (0.016)	0.010 (0.008)
Scale aggregate score	0.007*** (0.002)	0.008*** (0.003)	0.014*** (0.003)	0.011*** (0.003)	0.004 (0.003)	0.002 (0.002)
Constant	0.351*** (0.071)	0.328*** (0.060)	0.084 (0.094)	0.366*** (0.101)	0.409*** (0.098)	0.783*** (0.062)
Observations	458	579	579	518	504	528
R-squared	0.324	0.210	0.227	0.102	0.226	0.167

Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Table 5D
Caregivers' beliefs partial correlation with the variety of
play activities offered by the caregiver

Variety of play activities offered by caregiver	(1)	(2)
Treatment	0.509*** (0.144)	0.527*** (0.141)
Household wealth index (0-1)	1.034*** (0.340)	0.976*** (0.318)
Household members	-0.086*** (0.031)	-0.093*** (0.029)
Caregiver's age	-0.025*** (0.008)	-0.022*** (0.008)
Caregiver's educ. attainment (years)	0.099*** (0.019)	0.099*** (0.018)
Caregiver is married (yes = 1)	-0.192 (0.120)	-0.202* (0.120)
Hours worked last week by caregiver	0.003 (0.003)	0.003 (0.003)
Number of children under caregiver's care	-0.343 (0.285)	-0.364 (0.281)
Caregiver first language is Spanish (yes = 1)	0.089 (0.139)	0.042 (0.141)
Caregiver is the mother (yes = 1)	-0.509 (0.335)	-0.484 (0.343)
Child's age (months)	-0.019* (0.011)	-0.018 (0.011)
Child is male (yes = 1)	0.099 (0.123)	0.053 (0.109)
Baseline score	0.243*** (0.037)	0.238*** (0.036)
Scale aggregate score	0.053** (0.025)	
Scale average score		0.297* (0.170)
Constant	1.073 (0.792)	1.375* (0.805)
Observations	579	614
R-squared	0.290	0.287

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5E

Treatment effects by household wealth and caregiver education on caregivers' beliefs regarding the importance of parent-child interactions for child development (aggregate score)

(a) Caregiver education				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.173 (0.349)	0.924*** (0.319)	0.314 (0.400)	-0.353 (0.539)
Difference w.r.t Q1		0.751 (0.486)	0.141 (0.516)	-0.525 (0.617)
(b) Household wealth				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	-0.313 (0.365)	0.483 (0.361)	-0.382 (0.429)	1.039** (0.399)
Difference w.r.t Q1		0.796* (0.438)	-0.069 (0.614)	1.353** (0.551)

Number of observations is 579 (C = 304, T = 275).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions included pre-treatment controls.

Table 5F

Treatment effects on caregivers' beliefs regarding the importance of parent-child interactions for child development (caregivers with Spanish as mother tongue)

	Obs.	Mean control group	Mean treatment group	Difference	Controlled difference ^(a)	d ^(b)
Average score (1-4)	393 (C=208; T=185)	2.984	3.040	0.056 (0.055)	0.046 (0.046)	0.122
Aggregate score (7-28)	379 (C=199; T=180)	20.965	21.3167	0.352 (0.375)	0.276 (0.331)	0.104

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

(a) Controlled difference is the estimated effect including pre-treatment controls.

(b) Effect size (d) calculated as the standardized controlled difference.

Table 5G

Treatment effects by household wealth and caregiver education on caregivers' beliefs regarding the importance of parent-child interactions for child development (caregivers with Spanish as mother tongue; average score)

(a) Caregiver education				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.007 (0.069)	0.177*** (0.058)	0.088 (0.072)	-0.053 (0.089)
Difference w.r.t Q1		0.170* (0.092)	0.081 (0.097)	-0.060 (0.098)
(b) Household wealth				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	-0.039 (0.058)	0.057 (0.072)	0.005 (0.105)	0.145** (0.066)
Difference w.r.t Q1		0.095 (0.081)	0.043 (0.124)	0.184* (0.091)

Number of observations is 393 (C = 208, T = 185).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions included pre-treatment controls.

Table 5H

Treatment effects by household wealth and caregiver education on caregivers' beliefs regarding the importance of parent-child interactions for child development (caregivers with Spanish as mother tongue; aggregate score)

(a) Caregiver education				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.042 (0.501)	1.175*** (0.411)	0.515 (0.518)	-0.416 (0.629)
Difference w.r.t Q1		1.133 (0.674)	0.473 (0.690)	-0.458 (0.702)
(b) Household wealth				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	-0.400 (0.425)	0.399 (0.521)	-0.077 (0.754)	1.012** (0.471)
Difference w.r.t Q1		0.799 (0.595)	0.322 (0.888)	1.412** (0.652)

Number of observations is 379 (C = 199, T = 180).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions included pre-treatment controls.

Appendix 5.2

Treatment effects on the variety of play material

Table 5I
Treatment effects on the variety of play material by household wealth and caregiver education quartiles


(a) Household wealth				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.937*** (0.291)	0.922*** (0.256)	0.850*** (0.283)	0.732*** (0.246)
Difference w.r.t Q1		-0.016 (0.355)	-0.088 (0.384)	-0.205 (0.371)
(b) Caregiver education				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.856*** (0.168)	0.962*** (0.301)	0.662** (0.310)	0.988*** (0.344)
Difference w.r.t Q1		0.106 (0.322)	-0.194 (0.363)	0.132 (0.358)

Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions include pre-treatment controls.

Figure 5A
Relationship between the variety of play material and caregivers' household wealth and education



Figures depict a locally weighted scatterplot smoothing of the variety of play material on household wealth (panel A) and caregiver education (panel B). In each graph, results are presented for the control and treatment groups. Therefore, treatment effects are given by the vertical distance between the two functions.

Appendix 5.3

Treatment effects on the variety of play activities according to the scores obtained in the caregiver beliefs scale

Table 5J

Treatment effects on the variety of play activities in different quartiles of the caregiver beliefs scale

(a) Average score				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.295 (0.197)	0.641*** (0.211)	0.254 (0.290)	0.791*** (0.191)
Difference w.r.t Q1		0.346 (0.213)	-0.041 (0.338)	0.496* (0.262)
(b) Aggregate score				
	First quartile (Q1)	Second quartile (Q2)	Third quartile (Q3)	Fourth quartile (Q4)
Treatment effect	0.328* (0.177)	0.658*** (0.233)	0.255 (0.289)	0.790*** (0.194)
Difference w.r.t Q1		0.330 (0.221)	-0.073 (0.321)	0.462* (0.254)


Number of observations is 614 (C = 320, T = 294).

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Treatment effects and differences estimated using an interaction between treatment status and membership to each wealth or caregiver education quartile. All regressions include pre-treatment controls.

Figure 5B

Relationship between the variety of play activities and the scores obtained in the caregiver beliefs scale



Figures depict a locally weighted scatterplot smoothing of the variety of play activities on the scores obtained in the caregiver beliefs scale (the average score in panel A and the aggregate score in panel B). In each graph, results are presented for the control and treatment groups. Therefore, treatment effects are given by the vertical distance between the two functions.

Chapter 6

Conclusions, policy implications, limitations and avenues for future research

6.1 Main conclusions and policy implications

Inequalities in children's cognitive development emerge at early age, persist as they grow older, and eventually translate into differences in several later-life outcomes related to their personal wellbeing. A robust body of knowledge and empirical evidence support the claim that these inequalities can be traced back to differences in the quantity and quality of parenting factors (or parental investments) present during early childhood.

Two related phenomena that have received less attention are: (i) the existence of other childhood influences relevant for the emergence of these gaps; and (ii) the extent to which parental investments are amenable to policy intervention. Motivated by this general line of enquiry and by the prominence of cognitive achievement gaps between urban and rural children in Peru, this thesis addressed four specific research questions. In particular:

1. What is the importance of school influences vis-à-vis that of the early learning and care environment when explaining inequalities in cognitive development between urban and rural children in Peru?
2. Can an intervention focused on parent-child interactions working at scale in rural Peru deliver an improvement in the quality of the home environment?
3. Can an intervention focused on parent-child interactions working at scale in rural Peru change parental behaviour so as to increase the amount of stimulation offered to the child?

4. What constraints faced by caregivers limit the effect of this intervention on caregiver engagement in educational play activities with the child?

Regarding the first research question, the analysis presented in Chapter 2 revealed that observable school influences occurring between ages 6 and 8 account for a significant share (around 35% and no less than 28%) of the difference in cognitive skill observed at age 8 between urban and rural children in Peru. The share attributable to differences in the early childhood environment is important but no larger than 50%. This analysis employed data from the Young Lives Study which comprises longitudinal information on cognitive test scores, family and school characteristics. The rich dataset and, especially, the decomposition strategy played an important role in offering a result less prone to biases than those reported so far in the literature.

Performing this kind of decomposition is problematic as one faces the risk of bias not only when estimating the individual effect of cognitive skill inputs, but also when assigning the contribution of specific variables to different categories (e.g. “school inputs” or “home inputs”). The thesis addressed these issues at length in Chapter 3. This discussion provided a stronger and more formal foundation for the decomposition strategy employed to answer the first research question. By linking the decomposition technique, its components and interpretation to explicit assumptions about the determination of skill and its inputs, the analysis presented in Chapter 3 also sought to highlight several key considerations pertinent to performing a linear decomposition of a cognitive achievement gap.

The second and third research questions were addressed in Chapter 4. For this, I used data on parenting practices collected from the control and treatment groups of families randomized for the evaluation of the home-visiting component of programme *Cuna Mas* in rural Peru. Results showed that a home-visiting intervention working at scale can produce a

positive effect on the quality of the home environment favourable for the child's cognitive and socio-emotional development. The effect is statistically significant ($p < 0.01$) and has a size ($d = 0.5$) comparable to that found for other interventions of much smaller scale and efficacy trials conducted in the developing world. Importantly, it was also possible to identify an increase in the variety of educational play activities offered to the child by the caregiver in addition to those occurring during home visits ($d = 0.3$; $p < 0.01$), which constitutes strong evidence of a shift in caregiver behaviour.

Finally, the fourth research question was addressed in Chapter 5. This chapter explored the constraints that limit the effect of a home-visiting intervention on caregivers' engagement in educational play activities with their children. In this regard, the literature places a strong emphasis on the transfer of "parenting skills"¹⁰⁸ to caregivers as a key determinant of programme success. This, I argued, carries the implicit assumption that the lack of these skills is the major constraint preventing parents from offering more stimulation opportunities to their children.

For this analysis I used data on parenting practices and caregiver beliefs regarding the importance of parent-child interactions for child development, collected from the same sample of families as the one considered for the exercise presented in Chapter 4. Results showed that treatment effects of the home-visiting component of programme *Cuna Mas* on caregiver behaviour are not significant among the poorest caregivers and exhibit a positive wealth gradient. In addition, no evidence was found suggesting that this heterogeneity is driven by an unequal transfer of parenting skills but, instead, evidence is consistent with caregiver beliefs conditioning their response to treatment. This suggests there are constraints besides the lack of parenting skills that determine the effect of the intervention and that

¹⁰⁸ Parenting skills are understood as knowledge regarding specific activities that can be carried out with children and the ability to perform them.

these seem related to the information available to caregivers regarding the importance of parent-child interactions for child development. This analysis served to draw attention to the resources that caregivers need to input to produce stimulation opportunities for their children, an exercise that can offer some useful policy lessons as will be explained below.

In fact, the four research questions addressed in this thesis were not only motivated by their potential contribution to knowledge but also by their policy relevance. Examining whether school influences play a significant role in shaping inequalities in cognitive skill formation and whether parenting practices can be affected by a home-visiting programme (and why this behavioural change can happen), has clear policy implications.

In this regard, the results of this thesis confirm that early childhood influences play a significant role in explaining developmental setbacks affecting children from disadvantaged families. In fact, the literature has already presented evidence indicating the importance of early parental investments and has discussed how policy action focused on early childhood faces no trade-off between equity and efficiency (Cunha and Heckman, 2007, 2008). Stimulation opportunities offered by caregivers are a significant component of these early-life influences and this analysis has also shown that these parenting practices can be responsive to policy action focused on transferring knowledge to caregivers about specific activities that can be performed with their children.

Knowing that the home-visiting model can induce a change in caregiver behaviour constitutes an importance piece of evidence for two reasons. First, it validates the “theory of change” linking policy action with child development. Second, it means that the positive shift in the child’s environment has the potential of being sustained through time. If this occurs, there will be increases in child outcomes driven not only by the improved early childhood environment and its effect on the productivity of later inputs, but also by a shift in

these later inputs themselves. Evidence of behavioural change, therefore, makes a stronger case for policy interventions aimed at enhancing early childhood environments through the action of parents.

The results of this thesis also indicate that greater involvement in educational play activities comes at a cost for caregivers. From the perspective of an economic model of rational behaviour, this can be related to the opportunity cost of childcare time. For caregivers to be willing to devote more time to childcare, they must understand there will be a positive outcome they value. Consistent with this, the evidence presented here indicates that caregivers' response to the transfer of parenting skills is conditioned by their beliefs regarding the importance of parent-child interactions for child development.

The specific policy recommendation that stems from this evidence is to promote the inclusion of specific content regarding the importance of parent-child interactions as part of the information communicated to caregivers during the home visits. This can be complemented with the explicit consideration of achieving a shift in parental beliefs as one of the objectives of this type of interventions. It is worth noticing that a shift in parental beliefs reinforces the possibility of observing a permanent change in caregiver behaviour, something that, as already mentioned, promises a stronger effect on child development.

Results presented and discussed in this thesis further indicate that early childhood is not the only period when policy action can make a difference in terms of closing developmental gaps. A significant share of the urban/rural gap in cognitive development of Peruvian children at eight years of age can be explained by differences in the characteristics of their schools and teachers. The characteristics of rural schools and teachers have a direct connection with policy action because nearly all the supply of educational services in rural

areas is public. This leads to a clear policy implication:¹⁰⁹ efforts devoted to the equalization of the characteristics of rural schools and teachers with those existing in urban areas can allow a significant reduction in the cognitive skill gap between urban and rural children by the time they reach grade 3.

The fact that schools play an important role in shaping developmental gaps must serve to mitigate the pessimistic view expressed by some authors about the possibility of closing these differences through policy action. This pessimistic view stems from the notion that early childhood is the primary window of opportunity for policy, combined with the relatively small effects delivered by early childhood programmes when compared to the size of developmental gaps observed in the developing world (Schady et al., 2014).

Focusing solely on closing the gap between urban and rural schools and teachers will likely face a larger efficiency-equity trade-off than equalizing early childhood environments.¹¹⁰ This, however, does not mean that policy should refrain from aiming at reducing disparities in school services. In fact, if we bring together the results discussed in this thesis, we find strong evidence in favour of combining early childhood interventions focused on parent-child interactions with policy action aimed at improving the characteristics of rural schools and teachers. If the literature has already shown that early childhood interventions deliver only small effects with respect to the size of development gaps, improving school services for children from disadvantaged families can contribute an additional equalizing influence.

¹⁰⁹ The focus on skill inputs (rather than on input determinants) is a positive feature of this analysis, which allows one to draw useful policy implications. This is because inputs, by definition, have a direct effect on skill. If the supply of inputs found to have a significant contribution to some form of inequality can be directly affected by policy action, then policy can have an immediate role in mitigating this inequality.

¹¹⁰ Efficiency losses stem from the relatively small return to public investment targeting a population with low skill attainment. As documented in the first chapter of this thesis, skill attainment exhibits a strong positive correlation with socioeconomic status.

6.2 Limitations and avenues for future research

The analysis has some limitations that are worth considering. This exercise is useful not only because it allows a better appraisal of the results and their policy implications, but also because it can serve to identify opportunities for future research.

There are four limitations worth mentioning regarding the decomposition exercise. The first is related to the assumption of linearity imposed on the production function of skill. This simplification has allowed one to relate the decomposition strategy to the Oaxaca-Blinder technique employed in the literature and to focus on the issue of building variable categories in accordance with the role of inputs and input determinants in the production of skill. This has come at the cost of ignoring potential complementarities between inputs and, especially, between past cognitive achievement and contemporaneous influences (what the literature refers to as “dynamic complementarity”; see Cunha and Heckman (2007)). Here I have tested interactions between past cognitive test scores and school quality variables and found no significant results. Future research could test a more comprehensive set of interactions. Notice that, in this case, the decomposition strategy could include a special category to account for the contribution of concurrent differences in the inputs involved in the interaction.

A second limitation is related to the fact that the data considered only captures the effect of school influences up to third grade in primary education. Finding a significant contribution related to the first three years of schooling is surely reassuring about the importance of school influences in general. However, it would be interesting to reassess the contributions to the cognitive achievement gap in later grades in order to determine if contemporaneous home inputs (such as the amount of time devoted to study outside school) gain relevance as children progress through the schooling system. When publicly available, the fourth round

of data of the Young Lives study can provide cognitive scores and relevant household characteristics for this analysis. This has to be complemented with a new round of information on schools which, unfortunately and to the best of my knowledge, is not among the data collection plans of the study for Peru.

The third limitation relates to the “partial equilibrium approach” inherently followed by decomposition methods (Fortin et al., 2011). In particular, the shares estimated here correspond to a counterfactual exercise in which inputs belonging to other categories are held constant.¹¹¹ The total effect of equalizing school and teacher characteristics between urban and rural schools will ultimately depend on whether or not a shift in these characteristics triggers a change in the quantity and quality of inputs provided at home and, if it does, on whether families compensate or reinforce these changes.

The fourth limitation discussed here relates to the fact that the analysis is focused on school and teacher characteristics as a whole.¹¹² This implies that the results cannot inform policy about a specific school or teacher characteristic to prioritize in the effort of equalizing school services between rural and urban areas. Inspection of individual coefficients, however, allows one to identify a number of indicators of school quality which warrant further research in terms of their potential influence in compounding the educational disadvantage of rural pupils, including those relating to school climate and organisation. In particular, examining the qualitative differences between schools that relate to differences on indicators such as the use of multi-grade teaching and the incidence of teacher-reported “problems in class”, for example, may improve understanding of the differences in learning outcomes between urban and rural schools in the Peruvian context.

¹¹¹ This is consistent with a decomposition strategy based on contributions that depend on production function parameters. Recall that production function parameters, as opposed to conditional demand function parameters, indicate the partial effect of changing an input of skill.

¹¹² The decomposition exercise is based on consistent estimation of the contribution of a subset of inputs and not on the identification of specific production function parameters. As discussed in Chapter 2, the latter requires stronger assumptions.

Exploration of other potentially significant gaps in cognitive development represents an interesting opportunity for future research. In fact, the decomposition strategy employed in chapters 2 and 3 could be applied to the difference in cognitive development between indigenous and non-indigenous children or to the gap between children classified according to other predetermined skill influence such as gender or parental education.

Working with axes of inequality other than the urban/rural divide is a particularly promising avenue to explore gaps in socio-emotional skills. The Young Lives study has collected information on self-esteem and self-efficacy for its Younger Cohort starting at age 8 (see Dercon and Sanchez (2011) for a description of these personality traits and the scales employed to measure them). This implies that a value added specification could not be applied but, nonetheless, a cumulative model could provide interesting results regarding the relative importance of early childhood and school influences. This analysis, however, should focus on a gap different from that related to the urban/rural divide since the differences in these socio-emotional skills between urban and rural children are not as significant as those present in the cognitive domain (1 standard deviation; $p < 0.01$). In fact, the urban/rural gap in self-esteem amounts to only 0.19 standard deviations and remains marginally significant ($p < 0.1$) while the gap in self-efficacy is less than 0.02 standard deviations and is not statistically significant.

Before concluding, it is possible to highlight three limitations in connection with the analysis of the effects of a home-visiting intervention on parental behaviour. All three represent promising opportunities for future research aimed at further elucidating the mechanisms behind caregivers' behavioural change.

First, it was not possible to collect reliable information on caregiver time-use patterns. Accurate data on caregiver time-use would provide clearer indication about the extent to

which the intervention is demanding caregivers to input more time to childcare duties. A time-use diary was tested during the pilot study. The exercise revealed it was extremely problematic to elicit, through an interview, detailed information about the activities performed by rural caregivers at particular time slots throughout the day, as they do not keep track of time routinely. As a result, the instrument was discarded to concentrate the available resources on collecting reliable information on parenting practices.

Research about the effects of interventions focused on parent-child interactions on caregivers' time-use patterns is still lacking and represents a valuable opportunity to deepen our understanding of the mechanisms behind this type of programmes. The experience of this study reveals that this effort will require collecting information through more costly instruments such as those involving direct observation.

The second limitation is related to the low internal consistency exhibited by the scale on caregiver beliefs. Future research interested on measuring the effect of early childhood interventions on caregiver beliefs regarding their role in their children's development can devote additional effort to improve this instrument. An interesting option is discussed in Cunha et al. (2013). In this recent exercise conducted in the United States, researchers asked caregivers about the youngest age in which they believe certain developmental milestones can be attained by a child. Importantly, this was done under different hypothetical scenarios of parental investments, so an expected elasticity of child development to parental investment could be elicited. Further piloting will be necessary to validate this type of instrument in contexts where caregivers exhibit very limited numeracy and literacy skills such as in the case of the target population of *Cuna Mas*.

Finally, the third limitation has to do with the fact that data on caregiver beliefs were only collected during the follow-up survey and the intervention design did not allow for random

assignment of the scores related to this variable.¹¹³ The lack of pre-treatment information on caregiver beliefs limits the analysis of this variable as a potential moderator of treatment.¹¹⁴ As indicated above, this analysis has found evidence consistent with caregiver beliefs conditioning caregiver response to the transfer of parenting skills. A more conclusive statistical analysis, however, was not possible as correlation between treatment status and the scores obtained in the beliefs scale can introduce bias in the results.

Even if pre and post-treatment data on beliefs is available, it is not possible to ensure an unbiased assessment of this variable as a mediator of the intervention.¹¹⁵ This is because the scores obtained in the beliefs scale have not been randomly assigned, so correlation between this variable and unobservable influences can introduce bias in the estimate of treatment effects.

Two research opportunities related to the evaluation of future home-visiting interventions emerge from this discussion. First, a reassessment of caregiver beliefs as a moderator of these interventions can be conducted if information on this variable is collected as part of the baseline study. Second, strong causal evidence on the role of this variable as a mediator can be produced by, for example, randomizing the delivery of information about the importance of parent-child interactions for child development within the treatment group. This type of “encouragement design” would allow one to use the delivery of information as an instrument for caregiver beliefs when assessing the role of this variable as a mediator of the home-visiting intervention.

¹¹³ The baseline study was performed before the idea of analysing caregiver beliefs was proposed and it was not possible to influence the evaluation design.

¹¹⁴ Moderators influence the causal relationship between treatment status and the outcome on interest. They indicate for whom or under what conditions the treatment works (Kraemer et al., 2002).

¹¹⁵ Mediators identify the mechanisms through which a treatment works (Kraemer et al., 2002). If the home-visit intervention affects caregiver beliefs and these beliefs, in turn, affect caregiver engagement in educational play activities, one can say that beliefs mediate the effect of treatment on caregiver behavior.

References

- Almond, D., & Currie, J. (2011). Human Capital Development before Age Five. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4b).
- Andrabi, T., Das, J., Khwaja, A. I., & Zajonc, T. (2011). Do Value-Added Estimates Add Value? Accounting for Learning Dynamics. *American Economic Journal: Applied Economics*, 3(July), 29–54.
- Angrist, J., & Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*: Princeton University Press.
- Arteaga, I., & Glewwe, P. (2014). *Achievement Gap between Indigenous and Non-Indigenous Children in Peru: An Analysis of Young Lives Survey Data*. Young Lives Working Paper 130.
- Attanasio, O., Cattan, S., Fitzsimons, E., Meghir, C., & Rubio-Codina, M. (2015). *Estimating the Production Function for Human Capital: Results from a Randomized Control Trial in Colombia*. NBER Working Paper No. 20965. National Bureau of Economic Research.
- Attanasio, O., Fitzsimons, E., Fernandez, C., Grantham-McGregor, S., Meguir, C., & Rubio-Codina, M. (2012). *Home Visiting in Colombia: Impacts of a Scalable Intervention*. Centre for the Evaluation of Development Policies, Institute for Fiscal Studies.
- Attanasio, O., Grantham-McGregor, S., Fernandez, C., Fitzsimons, E., Rubio-Codina, M., & Meguir, C. (2013). Enriching the home environment of low-income families in Colombia: strategy to promote child development at scale. *Early Childhood Matters*, 120.
- Baker-Henningham, H., & Lopez-Boo, F. (2010). Early Childhood Stimulation Interventions in Developing Countries: A Comprehensive Literature Review. *IZA Discussion Paper Series*(5282).
- Baker-Henningham, H., Powell, C., Walker, S., & Grantham-McGregor, S. (2005). The effect of early stimulation on maternal depression: A cluster randomized controlled trial. *Archives of Disease in Childhood*, 90, 1230-1234.
- Baker, D., Goesling, B., & LeTendre, G. (2002). Socioeconomic Status, School Quality, and National Economic Development: A Cross-National Analysis of the “Heyneman-Loxley Effect” on Mathematics and Science Achievement. *Comparative Education Review*, 46, 291-312.
- Barrera-Orsorio, F., Garcia-Moreno, V., Patrinos, H. A., & Porta, E. (2011). *Using the Oaxaca-Blinder Decomposition Technique to Analyze Learning Outcomes Changes over Time: An Application to Indonesia's Results in PISA Mathematics*. Policy Research Working Paper 5584. The World Bank.
- Beltran, A., & Seinfeld, J. (2012). *La Trampa Educativa en el Peru: Cuando la Educacion Llega a Muchos pero sirve a Pocos*. Peru: Universidad del Pacifico.

- Bentley, M., Vazir, S., & Engle, P. (2010). A home-based educational intervention to caregivers in south India to improve complementary feeding and responsive feeding, and psychosocial stimulation increases dietary intake, growth and development of infants. *Journal of the Federation of American Societies for Experimental Biology*, 24.
- Biewen, M. (2012). *Additive Decompositions with Interaction Effects*. IZA DP No. 6730. Institute for the Study of Labor.
- Blinder, A. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 8(4), 436-455.
- Bornstein, M. H., & Putnick, D. L. (2012). Cognitive and Socioemotional Caregiving in Developing Countries. *Child Development*, 83(1), 46-61.
- Boyden, J., & Dercon, S. (2012). Child Development and Economic Development: Lessons and Future Challenges. *Young Lives working paper*.
- Bronfenbrenner, U. (1974). *A report on longitudinal evaluations of preschool programs: Vol. II. Is early intervention effective?* Washington DC: U.S. Department of Health, Education, and Welfare.
- Bronfenbrenner, U., & Morris, P. (2006). The Bioecological Model of Human Development. In R. M. Lerner & W. Damon (Eds.), *Handbook of Child Psychology: Theoretical Models of Human Development*. Hoboken NJ: Wiley.
- Burger, R. (2011). School effectiveness in Zambia: The origins of differences between rural and urban outcomes. *Development Southern Africa*, 28(2), 157-176.
- Burman, E. (1994). *Deconstructing Developmental Psychology*. London: Routledge.
- Caldwell, B., & Bradley, R. (1984). *Home observation for measurement of the environment*. Little Rock: University of Arkansas.
- Castro, J. F., Yamada, G., & Arias, O. (2012). *Higher Education Decisions in Peru: On the Role of Financial Constraints, Skills and Family Background*. CIUP Working paper DD/11/14. Centro de Investigación de la Universidad del Pacifico.
- Cohen, J. (1992). A Power Primer. *Psychological Bulletin*, 112, 155-159.
- Coleman, J., Campbell, E., Hobson, C., McPartland, J., Mood, A., Weinfeld, F., & York, R. (1966). *Equality of Educational Opportunity*. Washington, D.C: Department of Health, Education, and Welfare.
- Cooper, P., Landman, M., Tomlinson, M., Molteno, C., Swartz, L., & Murray, L. (2002). Impact of a mother-infant intervention in an indigent peri-urban South African context: Pilot study. *British Journal of Psychiatry*, 180, 76-81.
- Cooper, P., Tomlinson, M., Swartz, L., Landman, M., Molteno, C., Stein, A., . . . Murray, L. (2009). Improving quality of mother-infant relationship and infant attachment in socioeconomically deprived community in South Africa: randomised controlled trial. *British Medical Journal*, 338, 1-8.

- Crouch, L., Gustafsson, M., & Lavado, P. (2009). Measuring educational inequality in South Africa and Peru. In D. Holsinger & W. J. Jacob (Eds.), *Inequality in Education*: Springer Netherlands.
- Cueto, S., Guerrero, G., Leon, J., Zapata, M., & Freire, S. (2013). ¿La cuna marca las oportunidades y el rendimiento educativo? Una mirada al caso peruano. *Documento de Investigacion GRADE*, 66.
- Cueto, S., Guerrero, G., Leon, J., Zapata, M., & Freire, S. (2014). The relationship between socioeconomic status at age one, opportunities to learn and achievement in mathematics in fourth grade in Peru. *Oxford Review of Education*, 40, 50-72.
- Cueto, S., & Leon, J. (2012). *Psychometric Characteristics of Cognitive Development and Achievement Instruments in Round 3 of Young Lives*. Young Lives Technical Note 25.
- Cueto, S., Leon, J., & Munoz, I. (2014). Education Opportunities and Learning Outcomes of Children in Peru: A Longitudinal Model. In M. Bourdillon & J. Boyden (Eds.), *Growing Up in Poverty: Findings from Young Lives*: Palgrave Macmillan.
- Cunha, F., Elo, I. T., & Culhane, J. F. (2013). *Eliciting Maternal Subjective Expectations about the Technology of Cognitive Skill Formation*. NBER Working Paper 19144.
- Cunha, F., & Heckman, J. (2007). The Technology of Skill Formation. *American Economic Review*, 97(2), 31-47.
- Cunha, F., & Heckman, J. (2008). Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation. *The Journal of Human Resources*, 43(4), 738-782.
- Cunha, F., Heckman, J., Lochner, L., & Masterov, D. (2006). Interpreting the Evidence on Life Cycle Skill Formation. In E. Hanushek & F. Welch (Eds.), *Handbook of the Economics of Education* (pp. 697–812). Amsterdam: North-Holland.
- Cunha, F., Heckman, J., & Schenach, S. (2010). Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Econometrica*, 78(3), 883-931.
- Cunha, F., Heckman, J., & Schennach, S. (2010). Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Econometrica*, 78(3), 883-931.
- Deming, D., Hastings, J., Kane, T., & Staiger, D. (2014). School Choice, School Quality, and Postsecondary Attainment. *American Economic Review*, 104(3), 991–1013.
- Dercon, S., & Sanchez, A. (2011). *Long-Term Implications of Under-Nutrition on Psychosocial Competencies: Evidence from Four Developing Countries*. Young Lives Working Paper No. 72.
- Diaz, J. J., Arias, O., & Tudela, D. V. (2013). *Does Perseverance Pay as Much as Being Smart?: The Returns to Cognitive and Non-cognitive Skills in urban Peru*. Working paper. GRADE.

- Duflo, E., Glennerster, R., & Kremer, M. (2007). Using Randomization in Development Economics Research: A Toolkit. In T. P. Schultz & J. Strauss (Eds.), *Handbook of Development Economics* (Vol. 4): Elsevier.
- Dunn, L., Padilla, E., Lugo, D., & Dunn, L. (1986). *Manual del Examinador para el Test de Vocabulario en Imágenes Peabody: Adaptación Hispanoamericana*. Minnesota: AGS.
- Eickmann, S., Lima, A., Guerra, M., Lima, M., Lira, P., Huttly, S., & Ashworth, A. (2003). Improved cognitive and motor development in a community-based intervention of psychosocial stimulation in northeast Brazil. *Developmental Medicine and Child Neurology*, 45, 536-541.
- Elder, T., Goddeeris, J., & Haider, S. (2010). Unexplained gaps and Oaxaca–Blinder decompositions. *Labour Economics*, 17, 284-290.
- Engle, P. L., Black, M. M., Behrman, J. R., Mello, M. C. d., Gertler, P. J., Kapiriri, L., . . . Group, I. C. D. S. (2007). Strategies to avoid the loss of developmental potential in more than 200 million children in the developing world. *The Lancet*, 369(9557), 229-240.
- Engle, P. L., Fernal, L. C., Alderman, H., Behrman, J., O'Gara, C., Yousafzai, A., . . . Group, G. C. D. S. (2011). Strategies for reducing inequalities and improving developmental outcomes for young children in low-income and middle-income countries. *The Lancet*, 378, 1339-1353.
- Escobal, J., Lanata, C., Madrid, S., Penny, M., Saavedra, J., Suárez, P., . . . S.Huttly. (2003). *Young Lives Preliminary Country Report: Peru*. Young Lives Country Reports. Young Lives Study - Oxford Department of International Development.
- Farah, M., & Hackman, D. (2012). SES, Childhood Experience, and the Neural Bases of Cognition. In V. Maholmes & R. King (Eds.), *The Oxford Handbook of Poverty and Child Development*: Oxford Library of Psychology.
- Fiorini, M., & Keane, M. (2014). How the Allocation of Children's Time Affects Cognitive and Noncognitive Development. *Journal of Labor Economics*, 32, 787-836.
- Flavell, J. (1982). On Cognitive Development. *Child Development*, 53, 1-10.
- Fortin, N., Lemieux, T., & Firpo, S. (2011). Decomposition Methods in Economics. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4A): Elsevier.
- Foster, M. (2002). How Economists Think about Family Resources and Child Development. *Child Development*, 73, 1904-1914.
- Gardner, J., Walker, S., Powell, C., & Grantham-McGregor, S. (2003). A randomized controlled trial of a home-visiting intervention on cognition and behavior in term low birth weight infants. *The Journal of Pediatrics*, 143, 634-639.
- Glewwe, P., Hanushek, E., Humpage, S., & Ravina, R. (2013). School Resources and Educational Outcomes in Developing Countries: A Review of the Literature from

- 1990 to 2010. In P. Glewwe (Ed.), *Education Policy in Developing Countries*. Chicago: The University of Chicago Press.
- Glewwe, P., & Miguel, E. (2008). The Impact of Child Health and Nutrition on Education in Less Developed Countries. In T. P. Schultz & J. A. Strauss (Eds.), *Handbook of Development Economics* (Vol. 4): Elsevier.
- Gliem, J., & Gliem, R. (2003). *Calculating, Interpreting, and Reporting Cronbach's Alpha Reliability Coefficient for Likert-Type Scales*. Midwest Research to Practice Conference in Adult, Continuing, and Community Education.
- Grantham-McGregor, S., Cheung, Y. B., Cueto, S., Glewwe, P., Richter, L., Strupp, B., & Group, I. C. D. S. (2007). Developmental potential in the first 5 years for children in developing countries. *The Lancet*, 369(9555), 60-70.
- Grantham-McGregor, S., Powell, C., Walker, S., Chang, S., & Flether, P. (1994). The long term follow up of severely malnourished children who participated in an intervention program. *Child Development*, 65, 428-439.
- Grantham-McGregor, S., Powell, C., Walker, S., & Himes, J. (1991). Nutritional Supplementation, Psychosocial Stimulation and Mental Development of Stunted Children: The Jamaican Study. *The Lancet*, 338, 1-5.
- Grantham-McGregor, S., Schofiel, W., & Haggard, D. (1989). Maternal-child interaction in survivors of severe malnutrition who received psychosocial stimulation. *European Journal of Clinical Nutrition*, 43, 45-52.
- Grantham-McGregor, S., Schofiel, W., & Harris, L. (1983). Effects of psychosocial stimulation on the mental development of severely malnourished children: An interim report. *Pediatrics*, 72, 239-243.
- Grantham-McGregor, S., Schofiel, W., & Powell, C. (1987). Development of severely malnourished children who received psychosocial stimulation: 6 year follow up. *Pediatrics*, 79, 247-254.
- Guarino, C., Reckase, M. D., & Wooldridge, J. M. (2012). *Can value-added measures of teacher performance be trusted?* IZA Discussion Papers 6602.
- Guerrero, G., Leon, J., Rosales, E., Zapata, M., Freire, S., Saldarriaga, V., & Cueto, S. (2012). *Young Lives School Survey in Peru: Design and Initial Findings*.
- Guryan, J., Hurst, E., & Kearney, M. (2008). Parental Education and Parental Time With Children. *NBER Working Paper*(13993).
- Hamadani, J., Huda, S., Khatun, F., & Grantham-McGregor, S. (2006). Psychosocial Stimulation Improves the Development of Undernourished Children in Rural Bangladesh. *The Journal of Nutrition*, 136, 2645-2652.
- Hanushek, E., & Luque, J. (2003). Efficiency and equity in schools around the world. *Economics of Education Review*, 22, 481-502.

- Heckman, J. (2006). Skill Formation and the Economics of Investing in Disadvantaged Children. *Science*, 312, 1900-1902.
- Heckman, J. (2007). The economics, technology, and neuroscience of human capability formation. *Proceedings of the National Academy of Science*, 104(33), 13250-13255.
- Helmers, C., & Patnam, M. (2011). The formation and evolution of childhood skill acquisition: Evidence from India. *Journal of Development Economics*, 95, 252-266.
- Hernandez-Zavala, M., Patrinos, H., Sakellariou, C., & Shapiro, J. (2006). *Quality of Schooling and Quality of Schools for Indigenous Students in Guatemala, Mexico and Peru*. WPS3982. World Bank.
- Heyneman, S., & Loxley, W. (1982). Influences on Academic Achievement Across High and Low Income Countries: A Re-Analysis of IEA Data. *Sociology of Education*, 55, 13-21.
- Heyneman, S., & Loxley, W. (1983). The Effect of Primary-School Quality on Academic Achievement Across Twenty-nine High- and Low-Income Countries. *American Journal of Sociology*, 88, 1162-1194.
- Hoff, E., Laursen, B., & Tardif, T. (2002). Socioeconomic Status and Parenting. In M. Bornstein (Ed.), *Handbook of Parenting (Vol 2) Biology and Ecology of Parenting*: Lawrence Erlbaum Associates.
- Huberman, H., & Mendelsohn, A. (2012). Preventive Interventions: Parenting and the Home Environment. In V. Maholmes & R. King (Eds.), *The Oxford Handbook of Poverty and Child Development*: Oxford University Press.
- INEI. (2015). *Evolucion de la Pobreza Monetaria 2009-2014 Informe Tecnico*. Instituto Nacional de Estadística e Informatica.
- Jann, B. (2008). The Blinder–Oaxaca decomposition for linear regression models. *The Stata Journal*, 8, 453-479.
- Janssens, W., & Rosemberg, C. (2011). *The impact of a home-visiting early childhood intervention in the Caribbean on cognitive and socio-emotional child development*. Amsterdam Institute for International Development. Amsterdam.
- Jones, F., & Kelly, J. (1984). Decomposing Differences Between Groups: A Cautionary Note of Measuring Discrimination. *Sociological Methods & Research*, 12, 323-343.
- Kane, T., McCaffrey, D., Miller, T., & Staiger, D. (2013). *Have we identified effective teachers? Validating measures of effective teaching using random assignment*. Research Paper. MET Project. Technical report. Bill & Melinda Gates Foundation.
- Kimmel, J., & Connelly, R. (2007). Mothers' Time Choices Caregiving, Leisure, Home Production, and Paid Work. *The Journal of Human Resources*, 42, 643-681.
- Kraemer, H., Wilson, G., Fairburn, C., & Agras, W. (2002). Mediators and Moderators of Treatment Effects in Randomized Clinical Trials. *Arch Gen Psychiatry*, 59(10), 877-883.

- Krieg, J., & Storer, P. (2006). How Much Do Students Matter? Applying the Oaxaca Decomposition to Explain Determinants of Adequate Yearly Progress. *Contemporary Economic Policy*, 24, 563-581.
- Leibowitz, A. (1974). Home Investments in Children. In T. W. Schultz (Ed.), *Economics of the Family: Marriage, Children, and Human Capital*: National Bureau of Economic Research.
- Lopez-Boo, F. (2009). *The Production Function of Cognitive Skills: Nutrition, Parental Inputs and Caste Test Gaps in India*. Young Lives Working Paper No. 55.
- Magwaza, A., & Edwards, S. (1991). An evaluation of an integrated parent-effectiveness training and children's enrichment program for disadvantaged families. *South African Journal of Psychology*, 21, 21-25.
- McEwan, P. (2004). The Indigenous Test Score Gap in Bolivia and Chile. *Economic Development and Cultural Change*, 53, 157-190.
- McEwan, P., & Marshall, J. (2004). Why Does Academic Achievement Vary Across Countries? Evidence from Cuba and Mexico. *Education Economics*, 12(3), 205-217.
- McEwan, P., & Trowbridge, M. (2007). The achievement of indigenous students in Guatemalan primary schools. *International Journal of Educational Development*, 27, 61-76.
- MIDIS. (2013). *Nota Metodologica para la Evaluacion de Impacto del Programa Cuna Mas* Ministerio de Desarrollo de Inclusion Social. Peru.
- Muralidharan, K., & Sundararaman, V. (2013). *The aggregate effect of school choice evidence from a two-stage experiment in India*. Working Paper 19441. National Bureau of Economic Research.
- Nahar, B., Hamadani, J., Ahmed, T., Tofail, F., Rahman, A., S., H., & Grantham-McGregor, S. (2009). Effects of psychosocial stimulation on growth and development of severely malnourished children in a nutrition unit in Bangladesh. *European Journal of Clinical Nutrition*, 63, 725-731.
- Nair, M., Philip, E., Jeyaseelan, L., George, B., Matthews, S., & Padma, K. (2009). Effect of child development centre model early stimulation among at-risk babies. A randomized controlled trial. *Indian Pediatrics*, 46, S20-S26.
- Neumark, D. (1988). Employers' Discriminatory Behavior and the Estimation of Wage Discrimination. *The Journal of Human Resources*, 23, 279-295.
- Nores, M., & Barnett, W. S. (2010). Benefits of early childhood interventions across the world: (Under) Investing in the very young. *Economics of Education Review*, 29, 271-282.
- Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 14(3), 693-709.

- OECD. (2013). *PISA 2012 Results in Focus: What 15-year-olds know and what they can do with what they know*: OECD Publishing.
- Outes-Leon, I., Porter, C., & Sánchez, A. (2011). *Early Nutrition and Cognition in Peru: A Within-Sibling Investigation*. IDB Working Paper Series No. IDB-WP-241. Inter-American Development Bank.
- Outes-Leon, I., Sanchez, A., & Molina, O. (2010). *Psychosocial Status and Cognitive Achievement in Peru*. Young Lives Working Paper No. 65.
- Paxson, C., & Schady, N. (2007a). Cognitive Development among Young Children in Ecuador The Roles of Wealth, Health and Parenting. *The Journal of Human Resources*, XLII(1), 49-84.
- Paxson, C., & Schady, N. (2007b). Cognitive Development among Young Children in Ecuador The Roles of Wealth, Health, and Parenting. *The Journal of Human Resources*, XLII(1), 49-84.
- Peaker, G. (1971). *The Plowden Children Four Years Later*. London National Foundation for Educational Research in England and Wales.
- PNCM. (2013). *Jugando Apreñado (Guía para el Acompañante Técnico)*. Servicio de Acompañamiento a Familias del Programa Nacional Cuna Mas (PNCM).
- Powell, C. (2004). *An evaluation of the Roving Caregivers Program of the Rural Family Support Organisation*. UNICEF. Jamaica.
- Powell, C., Baker-Henningham, H., Walker, S., Gernay, J., & Grantham-McGregor, S. (2004). Feasibility of integrating early stimulation into primary care for undernourished Jamaican children: cluster randomised controlled trial. *British Medical Journal*, 329, 1-4.
- Powell, C., & Grantham-McGregor, S. (1989). Home visiting of varying frequency and child development. *Pediatrics*, 84, 157-164.
- Pungello, E., & Gardner-Neblett, N. (2012). Family Factors, Childcare Quality and Cognitive Outcomes. In V. Maholmes & R. King (Eds.), *The Oxford Handbook of Poverty and Child Development*: Oxford University Press.
- Ramos, R., Duque, J. C., & Nieto, S. (2012). *Decomposing the Rural-Urban Differential in Student Achievement in Colombia Using PISA Microdata*. IZA DP No. 6515. Institute for the Study of Labor.
- Rosenzweig, M., & Schultz, P. (1983). Estimating a Household Production Function: Heterogeneity, the Demand for Health Inputs, and Their Effects on Birth Weight. *Journal of Political Economy*, 91(5), 723-746.
- Schady, N., Behrman, J., Araujo, M. C., Azuero, R., Bernal, R., Bravo, D., . . . Vakis, R. (2014). *Wealth Gradients in Early Childhood Cognitive Development in Five Latin American Countries*. IDB-WP-482. Inter-American Development Bank.

- Singh, A. (2015). Private school effects in urban and rural India: Panel estimates at primary and secondary school ages. *Journal of Development Economics*, 113, 16-32.
- Tiebout, C. (1956). A Pure Theory of Local Public Expenditures. *Journal of Political Economy*, 64, 416-424.
- Todd, P., & Wolpin, K. (2003). On the Specification and Estimation of the Production Function for Cognitive Achievement. *The Economic Journal*, 113(February), F3 - F33.
- Todd, P., & Wolpin, K. (2007). The Production of Cognitive Achievement in Children: Home, School, and Racial Test Score Gaps. *Journal of Human Capital*, 1(1), 91-136.
- Totsika, V., & Sylva, K. (2004). The Home Observation for Measurement of the Environment Revisited. *Child and Adolescent Mental Health*, 9, 25-35.
- UNICEF. (2011). *Multiple Indicator Cluster Survey (MICS) - Questionnaire for Children Under Five*. United Nations Children's Fund.
- Walker, S. (2011). Promoting Equity through Early Child Development Interventions for Children from Birth through Three Years of Age. In H. Alderman (Ed.), *No Small Matter The Impact of Poverty, Shocks and Human Capital Investments in Early Childhood Development*. Washington, DC: The World Bank.
- Walker, S., Chang, S., Powell, C., & Grantham-McGregor, S. (2004). Psychosocial Intervention Improves the Development of Term Low-Birth-Weight Infants. *Journal of Nutrition*, 134, 1417-1423.
- Walker, S., Chang, S., Younger, N., & Grantham-McGregor, S. (2010). The effect of psychosocial stimulation on cognition and behaviour at 6 years in a cohort of term, low-birthweight Jamaican children. *Developmental Medicine and Child Neurology*, 52, e148-e154.
- Walker, S., Wachs, T., Gardner, J. M., Lozoff, B., Wasserman, G., Pollitt, E., . . . Group, I. C. D. S. (2007). Child development: risk factors for adverse outcomes in developing countries. *The Lancet*, Volume 369(9556), 145-157.
- WB. (2007). *Toward high-Quality Education in Peru: Standards, Accountability and Capacity Building*: World Bank.
- WB. (2014). World Development Indicators. Retrieved 24/11/2014, from World Bank <http://data.worldbank.org/indicator>
- Winsborough, H., & Dickenson, P. (1971). *Components of negro-white income differences*. Proceedings of the American Statistical Association, Social Statistics Section: 6-8.
- Yaqub, S. (2002). "Poor Children Grow into Poor Adults": Harmful Mechanisms or Over-Deterministic Theory? *Journal of International Development*, 14, 1081-1093.
- Zhang, L., & Lee, K. A. (2011). Decomposing achievement gaps among OECD countries. *Asia Pacific Education Review*, 12, 463-174.