NBER WORKING PAPER SERIES

ESTIMATING THE PRODUCTION FUNCTION FOR HUMAN CAPITAL: RESULTS FROM A RANDOMIZED CONTROL TRIAL IN COLOMBIA

Orazio Attanasio
Sarah Cattan
Emla Fitzsimons
Costas Meghir
Marta Rubio-Codina

Working Paper 20965
http://www.nber.org/papers/w20965

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2015

We thank participants at the NBER Summer Institute, Barcelona GSE Summer Forum and Montreal CIREQ Applied Microeconomics on Fertility and Child Development and seminars at Stanford University, University of Chicago, Oxford University, Cornell University, Bristol University and the Institute for Fiscal Studies for their comments. We are grateful to the Economic and Social Research Council (Grant ES/G015953/1), the Inter-American Development Bank, the International Growth Centre, and the World Bank for funding the intervention and data collection. Some of this research was financed by the European Research Council's Advanced Grant 249612 and by the Grand Challenges Canada Prime Award 0072-03 (sub-award reference number 560450). Sarah Cattan gratefully acknowledges financial assistance from the British Academy Postdoctoral Fellowship pf140104, as well as from the European Research Council's Grant Agreement No. 240910. Costas Meghir thanks the Cowles foundation and the ISPS at Yale for financial assistance. All errors are the responsibility of the authors. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w20965.ack

Estimating the Production Function for Human Capital: Results from a Randomized Control Trial in Colombia
Orazio Attanasio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina
NBER Working Paper No. 20965
February 2015
JEL No. H31,I24,I25,I28,I3,J38,O1,O15,O54

ABSTRACT

We examine the channels through which a randomized early childhood intervention in Colombia led to significant gains in cognitive and socio-emotional skills among a sample of disadvantaged children. We estimate production functions for cognitive and socio-emotional skills as a function of maternal skills and child's past skills, as well as material and time investments that are treated as endogenous. The effects of the program can be fully explained by increases in parental investments, which have strong effects on outcomes and are complementary to both maternal skills and child's past skills.

Orazio Attanasio
Department of Economics
University College London
Gower Street
London WC1E 6BT
UNITED KINGDOM
and NBER
o.attanasio@ucl.ac.uk

Sarah Cattan
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
sarah_c@ifs.org.uk

Emla Fitzsimons
Institute for Fiscal Studies
7 Ridgmount St
London, WC1E 7AE
UK
e.fitzsimons@ioe.ac.uk

Costas Meghir
Department of Economics
Yale University
37 Hillhouse Avenue
New Haven, CT 06511
and IZA
and also NBER
c.meghir@yale.edu

Marta Rubio-Codina
IFS
7 Ridgmount Street
London WC1E 7AE
United Kingdom
marta_r@ifs.org.uk

An online appendix is available at:
http://www.nber.org/data-appendix/w20965

1 Introduction

The first five years of life lay the basis for lifelong outcomes (Almond and Currie, 2011). Due to rapid brain development and its malleability during the early years (Knudsen, 2004; Knudsen et al., 2006), investments during this period may play a crucial role in the process of human capital accumulation. At this time, many children are, however, exposed to risk factors such as poverty, malnutrition and non-stimulating home environments preventing them from reaching their full potential and perpetuating poverty, particularly in developing countries (McGregor et al., 2007).

There is increasing evidence that early childhood interventions can alleviate the consequences of these detrimental factors in a long-lasting fashion. Examples include the Jamaica study (Gertler et al., 2014; Grantham-McGregor et al., 1991; Walker et al., 2011) and the Perry Preschool program (Heckman et al., 2010). However, less is known about the behavioural mechanisms through which these interventions affect children and their families.

This paper aims to examine the channels through which an early years intervention in Colombia led to significant gains in cognitive and socio-emotional skills among a sample of disadvantaged children. The intervention we study was a randomized control trial that targeted children between 12 and 24 months old, for a period of 18 months, in families who are beneficiaries of the conditional cash transfer program in Colombia (*Familias en Acción*). Its structure mirrored that of the Jamaica intervention in that it included a psychosocial stimulation component and a micro-nutrient supplementation component. The psychosocial stimulation program, which we focus on in this paper, provided weekly home visits to mothers of children, for a period of 18 months, with the aim of improving parenting practices in the early years and beyond.

The short-term impact evaluation of the intervention reveals that psychosocial stimulation had significant positive effects on the language and cognitive development of children who received the home visits (Attanasio et al., 2014). In what follows, we reproduce these results and also show impacts of the intervention in other dimensions, such as indicators of socio-emotional development. However, these results could have been generated by a number of different mechanisms. In addition to the weekly, one-hour home visit during which the child and their mother interacted with the home visitor, the intervention could have altered parental investment behavior by making them aware of the importance of early investments and informing them about parenting practices that enhance the child's learning at home.

In order to shed light on the mechanisms through which the stimulation program affected child development, we estimate parents' investment functions and the production functions for cognitive and socio-emotional skills. Following Cunha, Heckman, and Schennach (2010), we model the accumulation of future skills as a process that is determined by the child's current stock of skills, parents' investments and parental human capital as well as (unobservable) shocks. This technology is non-linear and allows the degree of substitutability between inputs to be determined from the data. We consider two types of investment (time and commodities) and allow parental choices to be endogenously determined by estimating investment functions that depend on resources and prices. This approach provides a natural framework to interpret and understand the potential channels through which the psychosocial stimulation component of the intervention could have boosted the skills of treated children. In particular, the intervention could have shifted the distribution of parental investments and/or changed the parameters of the production function, for example by making parents more productive or effective.

To estimate the production functions for cognitive and socio-emotional skills, we use data we collected both before and after the intervention in Colombia. The data contain very rich measures of child development, maternal skills and parental investments. Even with such rich data, estimating the parameters governing the skill formation process remains challenging for two reasons. First, inputs and outputs are likely to be measured with error. Second, inputs, especially investments, can be endogenous if parental investment decisions respond to shocks or inputs that are unobserved to the econometrician. To deal with the measurement error issue, we use dynamic latent factor models as Cunha, Heckman, and Schennach (2010). The endogeneity of investment is taken into account by implementing a control function approach, as in Attanasio, Meghir, and Nix (2015), whose estimation procedure we adopt here. The exclusion restrictions needed for identification are justified by the economics of the model.

Our estimates of the production function reveal a series of interesting and important findings. First, in line with the existing literature, we find strong evidence that a child's current stock of skills fosters the development of future skills.[1] Second, and also in line with the existing literature, we find that current skills, parental investments and maternal human capital are complementary in the production of future skills. Parental investments matter greatly for the accumulation of skills. In particular, material investments seem to matter more for cognitive skills, while time investments seem to matter more for socio-emotional skills. Our results indicate that the parameters that determine the productivity of investment greatly depend on whether investment is considered endogenous. Ignoring the fact that parents choose investment leads to a large downward bias of the estimated productivity of investment in the production functions, therefore indicating that parents use investment to compensate for negative shocks. Interestingly, this result is also obtained by Cunha, Heckman, and Schennach (2010) and Attanasio, Meghir, and Nix (2015), yet in very different contexts.[2]

With respect to the mechanisms through which the intervention operated, we find that the intervention significantly increased parental investments among treated families compared to non-treated ones. At the same time, there are no significant shifts in the parameters of the production function induced by the intervention. These two findings mean that the gains in cognitive and socio-emotional skills among children who received the intervention are fully

[1] These features of the technology of skill formation are often referred to as *self-productivity* and *cross-productivity* (Cunha et al., 2006).

[2] The former use the Children of the National Longitudinal Survey of Youth 1979, a longitudinal panel following the children of a representative sample of women born between 1956 and 1964 in the US. The latter use the Young Lives Survey for India, a longitudinal survey following the lives of children in two age-groups: a Younger Cohort of 2,000 children who were aged between 6 and 18 months when Round 1 of the survey was carried out in 2002, and an Older Cohort of 1,000 children then aged between 7.5 and 8.5 years. The survey was carried out again in late 2006 and in 2009 (when the younger children were about 8 the same age as the Older Cohort when the research started in 2002).

explained by the shift in investments.

Our findings make important contributions to the literature on human capital development, especially during the early years. To the best of our knowledge, our paper is the first to estimate the technology of skill formation in the first three years of life and to quantify the size of the dynamic complementarities between different domains of human development at such young ages. Characterising the production function at various ages is key for the identification of critical periods that are important to target for the development of particular skills. Our paper and that of Attanasio, Meghir, and Nix (2015) are the first to estimate non-linear production functions in a developing country context.[3] Helmers and Patnam (2011) estimate production functions with Indian data, but they rely on a linear technology, which implies that inputs are perfect substitutes for each other. Our results show that this assumption is strongly rejected by the data and that accounting for complementarities between inputs is of high importance. In this regard, our results are strikingly consistent with those of Cunha, Heckman, and Schennach (2010) and Attanasio, Meghir, and Nix (2015), neither of whom can reject that the technology of skill formation is Cobb-Douglas. Finally, we are the first to account for multiple types of investments in children. We establish that distinguishing between material and time investments is important for our understanding of skill formation in the early years. These findings have important implications for the design of future interventions.

While there is a vast literature evaluating the impact of early childhood

[3]Attanasio, Meghir, and Nix (2015) estimate nonlinear production functions for cognition and health in India for children from 5-15, using the Young Lives Survey.

interventions on child development, our paper innovates by complementing the information obtained from a randomized controlled trial of a specific intervention with a completely specified model of skill formation and parental investment in order to understand the mechanisms behind the observed impacts. In this sense, our paper shares the motivation of Heckman, Pinto, and Savelyev (2013), who are interested in the channels through which the Perry Pre-School Program produced gains in adult outcomes. Our focus and methodology, however, are different: Heckman, Pinto, and Savelyev (2013) perform a mediation analysis that decomposes linearly the treatment effects on adult outcomes into components attributable to early changes in different personality traits. We use a structural model in which parents make investment choices and human capital accumulates according to a completely specified production function to interpret and explain the impacts induced by a successful intervention. We explicitly test alternative and specific hypothesis about the origin of the impacts. Despite these differences, along with Heckman, Pinto, and Savelyev (2013) and a few other papers (Attanasio, Meghir, and Santiago, 2012; Duflo, Hanna, and Ryan, 2012; Todd and Wolpin, 2006), our paper illustrates how data from randomized trials can be profitably combined with behavioral models to go beyond the estimation of experimentally induced treatment effects and interpret the mechanisms underlying them, a crucial step for policy analysis.

The paper proceeds as follows. Section 2 describes the intervention and the data collected pre- and post-intervention and summarizes the short-term impacts of the intervention. Section 3 presents the theoretical framework we use and discusses its identification. Section 4 describes our estimation strategy. Section 5 presents the estimates of the model and discusses their implications for our understanding of the intervention. Section 6 concludes.

2 The intervention, its evaluation and its impacts

Although some influential studies have shown that well-designed and well-targeted interventions can achieve spectacular results that are sustained over long periods of time, a key challenge remains in the design of interventions that can be deployed on a large scale at reasonable cost whilst at the same time maintaining the quality that underlies the observed impacts. In this study, we use data from the evaluation of an intervention that was designed as an effectiveness rather than an efficacy trial as it was deployed on a relatively large scale and was delivered by local people. In this section, we give some details on the intervention and its evaluation design.

2.1 The intervention design

The integrated early childhood program analyzed in this paper was targeted at children aged between 12 and 24 months living in families receiving the Colombian CCT program (*Familias en Acción*), which targets the poorest 20% of households in the country. The intervention lasted 18 months, starting in early 2010. Appendix A contains a detailed description of the program's design, implementation and delivery. Here we summarize the key aspects.

The program was implemented in semi-urban municipalities in three regions of central Colombia, covering an area three times the size of England.

It had two components: psychological stimulation and micronutrient supplementation. In this paper, we focus on the psychosocial stimulation arm of the program. The stimulation curriculum was based on the Jamaican home visiting model, which obtained positive short- and long-term effects (Gertler et al., 2014; Grantham-McGregor et al., 1991; Walker et al., 2006). The protocols designed by Grantham-McGregor et al. (1991) for Jamaica were adapted to be culturally appropriate for Colombia. The aims of the home visits were to improve the quality of maternal-child interactions and to assist mothers to participate in developmentally-appropriate learning activities, centered around daily routines and using household resources.

Two key innovations vis-a-vis the Jamaican intervention were made so as to incorporate scalability and sustainability. The first was that the intervention was implemented on a much larger scale than in Jamaica. The second was that home visitors were drawn from a network of local women, generated by the administrative set-up of the CCT program. *Familias en Acción*, has, as many other CCTs, an important social function. Every 50-60 beneficiaries elect a representative who is in charge of organising social activities and who acts as mediators between them and the program administrators. These women, known as *Madre Líderes* (MLs), are beneficiaries of the program themselves and are typically more entrepreneurial and proactive than the average beneficiary. These characteristics marked them out as potentially effective home visitors.

The intervention was evaluated through a clustered randomized control trial. Within each of the three large regions, 32 municipalities were randomly assigned to one of 4 groups: (i) psychosocial stimulation, (ii) micronutrient supplementation, (iii) both, and (iv) control. Assignment to treatment was via cluster-level randomization. In each municipality, 3 MLs were selected and the children of the beneficiary households represented by each of these MLs and aged 12-24 months, were recruited to the study. Therefore, there was a total of 1,429 children living in 96 towns in central Colombia.

We conducted a baseline survey before the intervention started and a follow-up survey when it ended 18 months later. The surveys took place in the household, and children's development was measured directly by psychologists in community centers. The household surveys contain information on a rich set of socio-economic and demographic characteristics as well as less standard variables such as children's food intake, pre-school participation, maternal verbal ability and mental health, and maternal knowledge and information, amongst other things. We also collected information on stimulation in the home as reported by the mother, using the UNICEF Family Care Indicators (FCI) (Frongillo, Sywulka, and Kariger, 2003). This instrument includes questions about the types and numbers of play materials around the house and about the types and frequency of play activities the child engages in with an adult aged 15 or more (most often the mother).

Children's cognitive, language and motor development were assessed using the Bayley Scales of Infant and Toddler Development III, administered directly in community centers (Bayley, 2006). Children's language development was also assessed through maternal report using a Spanish adaptation of the short version of the MacArthur-Bates Communicative Development Inventory (Jackson-Maldonado, Marchman, and Fernald, 2012). Children's socio-emotional development was also measured through maternal report using the Bates' Infant Characteristics Questionnaire (Bates, Freeland, and Lounsbury, 1979) and the Early Children's Behavior Questionnaire (Putnam, Gartstein, and Rothbart, 2006). All of these tests were administered both pre- and post-intervention (using age-appropriate items), with the exception of the Early Children's Behavior Questionnaire which was only administered at follow-up. We describe them at length in Appendix B.

2.2 The short-term impacts of the intervention

2.2.1 Impacts on child development

The top panel of Table 1 summarizes the short-term impact of the intervention on measures of cognitive and socio-emotional development, some of which are reported in Attanasio et al. (2014).[4] The short-term impact evaluation of the home visits showed an increase of 0.24 of a standard deviation (SD) in cognitive development and an increase of 0.17 SD in receptive language, assessed using the Bayley Scales of Infant and Toddler Development (Bayley-III).[5]

The lower panel of the table also shows that the intervention led to an im-

[4] At baseline, we administered the Bayley-III to 1,420 children and the survey to 1,429 households (Figure 1). We excluded from analyses two children who scored less than three standard deviations below the mean on the Bayley-III cognitive subscale. The attrition rate between baseline and follow-up for the Bayley-III sample was approximately 10.62% (n=153) across treatment arms: 36 (10.00%) of the children from the stimulation arm were not measured at follow-up, 47 (13.06%) from the supplementation arm, 39 (10.83%) from the combined arm and 31 (8.61%) from the control arm. The difference in loss among the groups was not statistically significant.

[5] These treatment effects are slightly different from those reported in Attanasio et al. (2014) because in this paper we estimate the impact of psychosocial stimulation by pooling the two groups that received it and the two groups that did not, while Attanasio et al. (2014) estimates the impact of each of the four arms of the intervention separately.

provement in some dimensions of socio-emotional development. In particular, it resulted in a 0.07 SD decrease in the dimension of the Bates scale measuring difficult behavior; none of the other three components of the Bates scale were significantly affected, however.

As discussed in greater length in Attanasio et al. (2014), no significant impact of micro-nutrient supplementation on any child developmental outcomes were found. As a result, in this paper, we focus on understanding the effect of the psychosocial stimulation program on cognitive, language and socio-emotional development.

2.2.2 Suggestive evidence of mechanisms

There are various mechanisms through which the psychosocial stimulation program could have been effective in improving children's cognitive, language and socio-emotional development. The one-hour weekly visit aimed at providing mothers with information on early childhood development and demonstrating to them various developmental play activities they could repeat with their child in between weekly home visits. The materials and toys used in the visit were left in the home for the week following the visit in order to promote increased interaction (both quality and quantity) between mother and child on an ongoing basis. This should have subsequently affected positively various aspects of the child's home environment, as well as the mother, whose self-esteem, mental health[6] and parenting activities might have improved.

[6]We tested for impacts of the intervention on the mother's mental health, years of education, IQ, vocabulary and maternal knowledge as measured by the Knowledge of Infant Development Inventory (KIDI) (see Appendix B for a detailed description of the scales we use to measure these dimensions). We did not detect any significant impacts on these dimensions

The lower panel of Table 1 summarizes some of the results from the short-term evaluation. These show large increases in intervention areas in the varieties of play materials and play activities in the home, as measured by the FCI. This is indicative that one mechanism through which home visits might have improved child development was by promoting parental investments in children. In order to test this hypothesis and assess the extent to which changes in parental investments contributed to the observed impacts of the intervention, we need a framework to understand the process of child development. We use a production function to model the relationship between inputs and the output of skill, which we describe below.

3 The accumulation of human capital in the early years: a theoretical framework

In the previous section, we reported some of the impacts that an early years intervention had both on children developmental outcomes and on parental behavior. These estimates were straightforward to obtain due to the presence of a cluster randomized control trial designed to evaluate the intervention. We now build a theoretical framework that can be used to interpret and understand these results.

In particular, we use a production function to describe the process through which the skills of children evolve between the beginning and the end of the intervention. We refer to the baseline period as t, when children were aged

of mothers' human capital and therefore only report impacts on parental investments as measured by the FCI in Table 1. We return to this issue later in the paper.

Table 1: Short-term impacts of psycho-social stimulation on cognition, language, and find motor development; child temperament; and parental investments

Instrument:	Bayley				MacArthur-Bates	
Item:	Cognitive	Language receptive	Language expressive	Fine motor	Vocabulary	Complex sentences
Treatment effect	0.244**	0.175**	0.0320	0.0713	0.0947	0.0606
	(0.0621)	(0.0647)	(0.0623)	(0.0617)	(0.0652)	(0.0563)
Observations	1,264	1,264	1,262	1,261	1,321	1,321

Instrument	Bates				Family Care Indicators (FCI)	
Item:	Unsociable	Difficult	Unadaptable	Unstoppable	Varieties of play materials	Varieties of play activities
Treatment	-0.0433	-0.0758+	0.0597	-0.0313	0.213**	0.273**
	(0.0549)	(0.0455)	(0.0615)	(0.0535)	(0.0637)	(0.0499)
Observations	1,326	1,326	1,326	1,326	1,326	1,326

Notes: The unit of observation is the child. Coefficients and standard errors (in parentheses) from a regression of the dependent variable measured at follow-up on the intervention variable (a treatment dummy for psychosocial stimulation, combining children receiving stimulation alone and children receiving both stimulation and micro-nutrient supplementation) controlling for: child's sex; baseline level of the outcome (except for MacArthur-Bates "Complex sentences", where we control for baseline number of words spoken because the item measuring "Complex sentences" was not measured at baseline); and tester dummies. Standard errors are adjusted for clustering at the municipality level. **, * and + indicate significance at 1, 5, and 10%. All scores have been internally standardized non-parametrically for age and are therefore expressed in standard deviations (see Appendix B for details about the measures and the standardization procedure).

between 12 to 24 months old, and to the post-intervention period as $t + 1$, when children were aged between 30 to 42 months old. Children's skills at time $t + 1$ are assumed to be a function of the vector of skills at t, of parental skills, of parental investments and of some shocks. Our first aim is to characterize such a function and estimate its parameters. We assume that parents choose investments in human capital, reflecting their taste, their resources and information about the current evolution of skills. Together with the production function we estimate an investment function. Finally, we explicitly allow for measurement error of all the relevant variables that enter the production function and are determined by it: child and parents' skills as well as parents' investment.

Within this framework, the intervention can affect the accumulation of skills through different channels. For example, the intervention can change the parameters of the production function or can change parents' investment behavior. To allow for these effects, we let some of the parameters of the production function and of the investment function depend on the (randomly allocated) intervention.

Because we only focus on the effect of the psychosocial stimulation program, we define the non-treated group ($d = 0$) as the group of children who did not receive the home visits (therefore including both the control group and the group who only received the micro-nutrients) and the treated group ($d = 1$) as the group of children who received the home visits (therefore including those who received only the home visits and those who received both the home visits and the micro-nutrients).

3.1 The production function for human capital

We consider a two-dimensional vector of skills, which includes cognitive and socio-emotional skills. In the baseline period, child i's skills are denoted $\theta_{i,t} = (\theta^C_{i,t}, \theta^S_{i,t})$, where $\theta^C_{i,t}$ and $\theta^S_{i,t}$ are cognitive and socio-emotional skills at t, respectively. At the end of the intervention, the child's skills are denoted $\theta_{i,t+1} = (\theta^C_{i,t+1}, \theta^S_{i,t+1})$.

Following Cunha, Heckman, and Schennach (2010), we assume that the stock of skills in period $t+1$ is determined by the baseline stock of the child's cognitive and socio-emotional skills $\theta_{i,t}$, the mother's cognitive and socio-emotional skills, denoted by $P^C_{i,t}$ and $P^S_{i,t}$ respectively, and the investments $I_{i,t}$ made by the parents between t and $t+1$.[7] We also allow for the effect of a variable $\eta^k_{i,t}$ that reflects unobserved shock or omitted inputs. As with skills, parental investments $I_{i,t}$ can be a multi-dimensional vector. Here, we distinguish between material and time investments, which we denote as $I^M_{i,t}$ and $I^T_{i,t}$ respectively.

For each skill, we assume the production function is of the Constant Elasticity of Substitution (CES) type, so we can write the technology of formation for skill k as follows:

$$\begin{aligned}\theta^k_{i,t+1} =& A^k_d[\gamma^k_{1,d}\theta^{C\,\rho_k}_{i,t} + \gamma^k_{2,d}\theta^{S\,\rho_k}_{i,t} + \gamma^k_{3,d}P^{C\rho_k}_{i,t} + \gamma^k_{4,d}P^{S\,\rho_k}_{i,t} \\ &+ \gamma^k_{5,d}I^{M\rho_k}_{i,t} + \gamma^k_{6,d}I^{T\,\rho_k}_{i,t}]^{\frac{1}{\rho_k}} e^{\eta^k_{i,t}} \qquad\qquad k \in \{C, S\}\end{aligned} \tag{1}$$

where A^k_d is a factor-neutral productivity parameter and $\rho_k \in (-\infty, 1]$ deter-

[7]Note that because the mother is the main caregiver in most families, we focus on her skills as those that are most likely to influence the child's development.

mines the elasticity of substitution, given by $1/(1-\rho_k)$, between the inputs affecting the accumulation of skill k. Under such parameterization, as $\rho_k \to -\infty$, the inputs become perfect complements. As $\rho_k \to 1$, the inputs become perfect substitutes. Notice that we let all the parameters of the production function, except the elasticity of substitution ρ_k, be a function of the intervention. This choice is dictated by our interpretation of how the intervention could have generated the impacts documented above.

First, the intervention could have changed the parameters of the production function that determine productivity. For example, by providing information about good parenting practices, the intervention could have increased the quality of the investments. In the framework above, this could be reflected by a shift in the factor-neutral productivity parameter A_d^k between the treated and the non-treated group or a shift in particular share parameters $\gamma_{j,d}^k$. Second, the intervention could have affected $I_{i,t}^M$ and $I_{i,t}^T$, the level of material and time investments that parents make, as suggested by the results presented in Table 1. For this reason we will let the parameters of the investment function, which we describe below, be a function of the intervention. In addition, the intervention could also have affected mothers' skills, for instance by improving self-esteem and reducing depression, or maternal knowledge as measured by the Knowledge of Infant Development Inventory (KIDI). Although we checked for these impacts, we did not detect any differences in measures of maternal human capital between control and treated after the intervention, so going forward we assume this mechanism away. From now on, we will therefore assume that mother's human capital is time-invariant and denote her cognitive and socio-emotional skills by P_i^C and P_i^S.

A few other features of the production function should be noted. First, all the parameters are specific to a particular skill, so the productivity parameter, the share parameters and the elasticity substitution can differ between the production function of the cognitive skills and that of the socio-emotional skills. Second, the CES functional form provides a great level of flexibility in that it allows the degree of substitutability between the various inputs of the production function to be determined by the data and to range from perfect substitutes to perfect complements. One well-known limitation of the CES functional form is that it imposes the same elasticity of substitution between any two inputs. This could, of course, be alleviated by estimating more general production functions, and in preliminary work we experimented with nested CES production functions. We could not reject the CES functional form however and so we maintain this functional form assumption throughout the application.

There are two main challenges to identifying and estimating the production functions outlined above. The first is related to the fact that parents choose investment in children, so that investments are likely to be correlated with the unobserved shock $\eta_{i,t}^k$. The second issue is that children's skills, as well as parental investments and maternal skills, could be measured with error in the data. We discuss how we tackle these two issues in the sub-sections below.

3.2 Accounting for the endogeneity of parental investments

The first issue that complicates the identification of the production function outlined above is the possibility that parental investments are endogenous, which would arise if $E(\eta_{i,t}^{k}|I_{i,t}) \neq 0$. There are two main reasons why parental investments might be correlated with the unobserved shock affecting the accumulation of human capital. First, parental investments might be correlated with omitted inputs in the production function of the child's skills. Second, parental investments might respond to unobserved, time-varying shocks in order to compensate or reinforce their effects on child development. Consider, for example, the case of a child who is suddenly affected by a negative shock, such as an illness, which is unobserved to the econometrician but perceived by the parents as delaying the child's development. As a result of this shock, parents might decide to invest in their child's development more than they would have otherwise. This parental response would create a negative correlation between parental investments and the unobserved shock $\eta_{i,t}^{k}$ affecting the development of skills of type k, which would lead the estimate of the effect of parental investments on future skill to be downward biased. Alternative assumptions about preferences and technologies (or technologies as perceived by the parents) can create different patterns of correlations between shocks and investment and, therefore, introduce different types of biases.

Endogeneity, of course, arises because parents choose investment in children to maximize some objective function taking into account the technology of human capital accumulation, the costs of investment and the resources available. In such a context, investment choices in any period will depend on initial conditions, on the shocks affecting the child and observed by the parent, on prices and on total resources. Rather than modeling investment choices jointly with the production function and making specific assumptions on taste (which would imply a specific functional form for the investment function), we estimate a reduced form equation that should be interpreted as an approximation of the investment function. We then use this approximation to implement a control function approach in the estimation of the production function for human capital and, therefore, control for the endogeneity of investment.

For identification, our approach requires that some variables that determine investment choices do not enter the production function directly. A natural candidate would be the intervention we described above, as it was allocated randomly across villages. However, as we want to test whether the intervention changed the parameters of the production function, we cannot use it as an exclusion restriction and need additional variables for valid identification. Moreover, as we model separately different forms of investment, the intervention alone would not be sufficient.

Following Cunha, Heckman, and Schennach (2010), we assume that, for each factor k, the error term $\eta_{i,t}^k$ can be decomposed into two components, $\pi_{i,t}^k$ and $\upsilon_{i,t}^k$. The production function can then be re-written in logs as:

$$\begin{aligned}\ln(\theta_{i,t+1}^k) = \;& \frac{1}{\rho_k}\ln[\gamma_{1,d}^k \theta_{i,t}^{C\,\rho_k} + \gamma_{2,d}^k \theta_{i,t}^{S\,\rho_k} + \gamma_{3,d}^k P_i^{C\rho_k} + \gamma_{4,d}^k P_i^{S\rho_k} + \gamma_{5,d}^k I_{i,t}^{M\,\rho_k} \\ & + \gamma_{6,d}^k I_{i,t}^{T\,\rho_k}] + \ln(A_d^k) + \delta^k \pi_{i,t}^k + \upsilon_{i,t}^k, \qquad k = \{C, S\}\end{aligned} \tag{2}$$

Both π_t^k and $\upsilon_{k,t}$ are assumed to be distributed independently across children. However, π_t^k is assumed to be realized before parents make investment choices and therefore can influence their choices, whereas $\upsilon^k_{i,t}$ is realized after parents make investment choices. The goal of the control function approach is to recover a consistent estimate of $\pi^k_{i,t}$, so that it can be controlled for when estimating the production function. This in turn requires estimating a model of investment.

As discussed above, we do not derive explicit investment functions from a complete structural model. Instead, we specify an approximation to non-linear investment functions as log-linear equations in initial conditions, maternal skills and the variables $Z_{i,t}$, representing resources and prices:

$$\begin{aligned}\ln(I^{\tau}_{i,t}) =& \lambda^{\tau}_{d,0} + \lambda^{\tau}_{d,1}\ln(\theta^{C}_{i,t}) + \lambda^{\tau}_{d,2}\ln(\theta^{N}_{i,t}) + \lambda^{\tau}_{d,3}\ln(P^{C}_{i}) + \lambda^{\tau}_{d,4}\ln(P^{S}_{i}) \\ &+ \lambda^{\tau}_{d,5}\ln(Z_{i,t}) + u^{\tau}_{i,t}, \qquad \tau = \{M, T\}\end{aligned} \tag{3}$$

where $u^{\tau}_{i,t}$ is a linear combination of $\pi^C_{i,t}$ and $\pi^S_{i,t}$. Note that all parameters of the investment functions, including the intercept, are allowed to vary between the treated and non-treated groups. This reflects the possibility, discussed above, that the intervention changed parental investment strategies.

Once the parameters of the investment functions are estimated, we recover $\hat{u}^T_{i,t}$ and $\hat{u}^M_{i,t}$, the estimated residuals from the investment equations (3), which we include as regressors when estimating the production functions:

$$\begin{aligned}\ln(\theta^{k}_{i,t+1}) =& \frac{1}{\rho_k}\ln[\gamma^{k}_{1,d}{\theta^{C}_{i,t}}^{\rho_k} + \gamma^{k}_{2,d}{\theta^{S}_{i,t}}^{\rho_k} + \gamma^{k}_{3,d}{P^{C}_{i,t}}^{\rho_k} + \gamma^{k}_{4,d}{P^{S}_{i,t}}^{\rho_k} + \gamma^{k}_{5,d}{I^{M}_{i,t}}^{\rho_k} \\ &+ \gamma^{k}_{6,d}{I^{T}_{i,t}}^{\rho_k}] + \ln(A^{k}_{d}) + \phi^{k}_{M}\hat{u}^{M}_{i,t} + \phi^{k}_{T}\hat{u}^{T}_{i,t} + \upsilon^{k}_{i,t}, \qquad k = C, N\end{aligned} \tag{4}$$

Identification of the parameters of the production function rests on the assumption that the disturbances $\pi_{i,t}^k$ and $\upsilon_{i,t}^k$ in equation (2) are independent of $Z_{i,t}$ and that there are at least as many exclusion restrictions – variables that affect the technology of skill formation only through the investment process – as there are endogenous variables. Economic theory suggests that variables that exogenously shift the household's resources might be valid exclusion restrictions, since they impact parental investment decisions through the budget constraint without entering directly the production function. In this spirit, we use average male and female wages in the child's village, household's wealth at baseline, and an indicator for whether the mother is married as variables that determine resources but do not enter the production function explicitly. These variables are valid exclusion restrictions insofar as, conditional on the child's skills at baseline and maternal human capital, they are orthogonal to π_t^k. We believe that this assumption is likely to hold, as we control for a multitude of child and parents' characteristics through the latent factors.

One of the variables we use as a determinant of investment that does not enter the production function is household wealth. One could argue that household wealth is endogenous to unobserved shocks affecting the child. Indeed, going back to our example above, it is possible that parents of a sick child decide to work more and increase their wealth in order to bolster the care they can provide him or her. Our strategy is less likely to suffer from this caveat, however, because we use baseline measures of wealth that should precede any shocks occurring between t and $t+1$. Moreover, excluding wealth from the investment functions does not change the point estimates but only affects precision.

3.3 Measurement of skills and investments

As described in section 2.2, the data contains multiple measures of the inputs and outputs of the production functions specified above. These measures are likely to proxy a lower-dimensional vector of skills and investment, but to do so with some error. In order to deal with this issue, we follow Cunha, Heckman, and Schennach (2010) in using latent factor models and we estimate the joint distribution of error-ridden latent factors measuring children's skills at baseline (θ_t^C and θ_t^S), children's skills at follow-up (θ_{t+1}^C and θ_{t+1}^S), mother's skills (P^C and P^S) and parental investments (I_t^M and I_t^T), where we keep the individual subscript i implicit for notational simplicity. Our specific estimation approach follows the procedure developed in Attanasio, Meghir, and Nix (2015) and is detailed in the next section.

Suppose we have $\mathcal{M}_{k,t}^1$ measures of child's skills of type k ($k \in \{C, S\}$) at time t. We also have $\mathcal{M}_k^2$ measures of maternal skills of type k ($k \in \{C, S\}$) and $\mathcal{M}_{\tau,t}^3$ measures of parental investments of type τ ($\tau = \{M, T\}$) at t.[8] Let $m_{k,t,j}^1$ denote the jth measure of child's skill of type k at t, $m_{k,j}^2$ the jth measure of mother's skill of type k, and $m_{\tau,t,j}^3$ the jth measure of parental investment of type τ at t. As is common in the psychometric literature, we assume a dedicated measurement system, that is one in which each measure only proxies one factor (Gorusch, 1983, 2003).[9] Assuming each measure is

[8]The measures of maternal skills are not indexed by time because we have assumed they are time-invariant.

[9]This assumption is not necessary for identification, but we choose to specify a dedicated measurement system so as to make the interpretation of the latent factors more transparent. As described in Appendix C, we find clear support in the data for such a system.

additively separable in the (log) of the latent factor it proxies,[10] we can write:

$$m^1_{k,t,j} = \mu^1_{k,t,j} + \alpha^1_{k,t,j} \ln \theta^k_t + \epsilon^1_{k,t,j} \quad (5)$$

$$m^2_{k,j} = \mu^2_{k,j} + \alpha^2_{k,j} \ln P^k + \epsilon^2_{k,j} \quad (6)$$

$$m^3_{\tau,t,j} = \mu^3_{\tau,t,j} + \alpha^3_{\tau,t,j} \ln I^\tau_t + \epsilon^3_{\tau,t,j} \quad (7)$$

where the terms $\mu^1_{k,t,j}, \mu^2_{k,j}$ and $\mu^3_{\tau,t,j}$ are intercepts, the terms $\alpha^1_{k,t,j}, \alpha^2_{k,j}$ and $\alpha^3_{\tau,t,j}$ are factor loadings, and the terms $\epsilon^1_{k,t,j}, \epsilon^2_{k,j}$ and $\epsilon^3_{\tau,t,j}$ are measurement errors. Note that the latent factors can be freely correlated with each other.

An important specificity of our application of latent factors models is that we consider an intervention and aim to capture its effect on the entire distribution of latent factors. To do so, we allow the joint distribution of the latent factors to be completely different between the two treatment states ($d = \{0, 1\}$). In contrast, we assume that the intercepts, factor loadings and measurement errors are invariant across states. These assumptions imply that any difference in the distribution of measures between the control and treated groups result from differences in the distribution of the latent factors and not from differences in the measurement system for those factors. As discussed in Heckman, Pinto, and Savelyev (2013), these assumptions are sufficient but not necessary for identification. We maintain them in our application because they restrict the number of free parameters and lead to improvements in efficiency. Additionally, because the treatment was randomized successfully, there is no reason to think that these parameters should vary across groups.

[10] We specify the measurement equation such that measures proxy the log of a latent factor so that latent factors only take positive values.

Because the latent factors are unobserved, identification of factor models requires normalizations to set their scale and location (Anderson and Rubin, 1956). We set the scale of the factors by setting the factor loading of the first measure of each latent factor to 1, that is: $\alpha^1_{k,t,1} = \alpha^2_{k,1} = \alpha^3_{\tau,t,1} = 1$, $\forall t, \tau = \{M, T\}$ and $k = \{C, S\}$. We set the location of the factors by fixing the mean of the latent factors to 0 in one group. Without loss of generality, we do so in the control group ($d = 0$) and allow the latent factor means of the treated group to be freely estimated.[11]

Under these normalizations, Heckman, Pinto, and Savelyev (2013) show that identification of the system is guaranteed as long as we have at least three measures dedicated to each factor under the assumptions that the measurement error is independent across measures and from the latent factors and that $E(\epsilon^1_{k,t,j}) = E(\epsilon^2_{k,j}) = E(\epsilon^3_{\tau,t,j}) = 0$ for $j \in \{1, \ldots, \mathcal{M}^1_{k,t}; 1, \ldots, \mathcal{M}^2_k; 1, \ldots, \mathcal{M}^3_{\tau,t}\}$, $\forall t, k \in \{C, S\}$ and $\tau \in \{M, T\}$.

Note that some of these assumptions could be relaxed (Carneiro, Hansen, and Heckman, 2003; Cunha and Heckman, 2008; Cunha, Heckman, and Schennach, 2010). For instance, the same measure could be allowed to load on several factors, as long as there are some dedicated measures. It would also be possible to allow measurement error to be correlated across measures of the same factor, as long as there was one measure whose measurement error was independent from the measurement error in other measures of the same factor.

[11] This normalization is innocuous because, without it, we would identify the difference in factor means between the treatment and control groups, which is exactly the object of interest.

4 Estimation of the model

The approach we use to estimate our model is described in detail in Attanasio, Meghir, and Nix (2015) and involves two main stages. In the first, we estimate the joint distribution of the latent factors; in the second, we estimate the parameters of the investment and production functions, using draws from the joint distribution of factors.

4.1 First stage: estimating the joint distribution of latent factors

As mentioned above, using the Kotlarski theorem and its extensions by Carneiro, Hansen, and Heckman (2003), Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010), the model is non-parametrically identified. For estimation however, we make some distributional assumptions for the distribution of latent factors and measurement error. In particular, we assume that the latent factors are distributed as a mixture of two joint log-normal distributions. The mixture of log normal distribution represents a flexible way to approximate a generic distribution. In principle, one could allow for a mixture of three or more log normal distributions for even greater flexibility, but in our application, we found the two-type mixture satisfactory. Note that it is important to allow for substantial deviation from normality, as such a functional form would imply linearity (or log-linearity in the case of log normality) of the conditional means and, in turn, restrict the elasticity of substitution among factors (for example, normality of the factors would imply that inputs are perfectly substitutable.)

For notational brevity, we denote from now on the vector of latent factors as

$\theta = (\theta^C_{t+1}, \theta^S_{t+1}, \theta^C_t, \theta^S_t, P^C, P^S, I^M_t, I^T_t)$. As mentioned above, we allow the joint distribution of θ to be fully-specific to each intervention group d ($d \in \{0, 1\}$) so their density function in group s can be written as:

$$p_d(\ln \theta) = \tau_d p_d(\ln \theta^A) + (1 - \tau_d) p_d(\ln \theta^B) \tag{8}$$

where $\ln \theta^A \sim N(\mu^A_d, \Sigma^A_d)$ and $\ln \theta^B \sim N(\mu^A_d, \Sigma^A_d)$ and τ_d is the mixture weight. In addition, we assume that the measurement errors are distributed as a joint normal distribution with means 0 and diagonal variance-covariance Σ_ϵ. Notice that an implication of the additive separability of the measurement equations (5) - (7), together with the assumption of log-normality of the factors and of normality of the additive measurement error, is that the joint distribution of measurements is given by a mixture of normals.

We first estimate the parameters of the joint distribution of measurements by maximum likelihood, using the EM algorithm. We then map these parameters into the parameters of the joint distribution of factors, the variances of measurement errors, the factor loadings and the intercepts and obtain estimates of these parameters by minimum distance. We report the relationships between the parameters of the distribution of measurements and those of the distribution of factors and of measurement errors in Appendix C. A more detail treatment of the approach is found in Attanasio, Meghir, and Nix (2015).

4.2 Second stage: estimating investment and production functions

Once the joint distribution of latent factors is estimated for each group d, we can estimate the investment functions and the production functions using draws from the estimated distributions as data. Given draws on all the latent factors, we first estimate the log-linear investment functions by ordinary least squares and construct the residuals $\hat{u}_t^\tau$ ($\tau \in \{M, T\}$) that serve as control functions. As mentioned above, we let the parameters of the investment function depend on the intervention, reflecting the fact that the intervention might have changed the way parents behave. Next, we estimate the parameters of the CES production functions by non-linear least squares, including the estimated residuals of the investment functions as additional regressors, as specified in equation (4).

To highlight the bias resulting from failing to account for the endogeneity of investments, we present results with and without control functions. We compute standard errors and confidence intervals using the bootstrap.[12]

4.3 Specification of the empirical model

In addition to the factors capturing the child's skills at baseline and mother's skills, in the production function, we include the number of children in the household (as measured at follow-up). This is to allow for the possibility that the presence of siblings affects child development, either because of spillover effects or by reducing the level of attention parents devote to each one of their children in multiple children households. We also include the number of children as a determinant of investment, which we suspect might depend negatively on the number of siblings. Since the number of children in the

[12]We draw $Q = 1000$ bootstrap samples of the original data, accounting for the fact that the data is clustered at the village level, and we apply the estimation procedure described above to each one of the pseudo-sample. For each of the parameters, we then compute the standard deviation of its distribution based on its $Q = 1000$ bootstrapped values, along with various percentiles to compute the corresponding confidence intervals.

household enters the production function directly, it does not enter our list of exclusion restrictions.

To compute our measure of household wealth, we add to the measurement system described above a set of measures from the baseline survey that proxy an additional latent factor measuring household wealth with error. These measures are described in full in Table 2 below and include indicators of whether the household owns its dwelling, along with various other assets (fridge, car, computer, etc.). The other exclusion restrictions (average male and female wages in the village and whether the mother is married), as well as the variable measuring the number of children in the household, are assumed not to have any measurement error.[13]

5 Results

We start by reporting the estimates of the measurement system, followed by the estimates of the investment and production functions. Finally, armed with these parameters, we assess how the model fits the data and how it helps us interpret the impact of the early years intervention we have studied.

[13]To estimate the joint distribution of all the data we need to estimate the investment and production functions, we therefore specify a measurement system that comprises of all the measures of child's skills, mother's skills, investments and household wealth, along with these four additional variables. Each of these four variables can be thought of being a function of a latent factor and a measurement error term (following the same structure as equation (5) for example), but in their case, the variance of measurement error is 0 and the associated factor loading is 1. With respect to male and female wages in the village, we took an average of male and females wages reported by members of the sample and in doing so rid the average of the measurement error possibly contained in individual observations.

Table 2: Measurement system and signal-to-noise ratio for each measures

Factor	Measures	% Signal Controls	% Signal Treated
Child's cognitive skills *(t+1)*	Bayley Cognitive	76%	77%
	Bayley Receptive Language	71%	72%
	Bayley Expressive Language	78%	79%
	Bayley Fine Motor	55%	57%
	Mac Arthur-Bates Vocabulary	55%	56%
	Mac Arthur-Bates Complex Sentences	38%	39%
Child's cognitive skills *(t)*	Bayley Cognitive*	74%	67%
	Bayley Receptive Language*	80%	74%
	Bayley Expressive Language*	80%	73%
	Bayley Fine Motor*	68%	60%
	Mac Arthur-Bates Vocabulary*	43%	35%
Child's socio-emotional skills *(t+1)*	Bates Difficult sub-scale (-)	69%	67%
	Bates Unsociable sub-scale (-)	21%	20%
	Bates Unstoppable sub-scale (-)	62%	60%
	Rothbart Inhibitory Control sub-scale	70%	68%
	Rothbart Attention sub-scale	25%	24%
Child's socio-emotional skills *(t)*	Bates Difficult factor* (-)	67%	72%
	Bates Unsociable factor* (-)	19%	23%
	Bates Unadaptable* (-)	34%	40%
	Bates Unstoppable* (-)	23%	28%
Wealth	Owns a fridge	39%	40%
	Owns a car	6%	6%
	Owns a computer	34%	35%
	Owns a blender	28%	29%
	Own a washing maching	8%	8%
	Owns dwelling	11%	11%
	Owns a radio	12%	13%
	Ows a TV	32%	33%

Factor	Measures	% Signal Controls	% Signal Treated
Material investments	Number of different play materials	96%	97%
	Number of colouring books	44%	46%
	Number of toys bought	65%	67%
	Number of toys that require movement	73%	75%
	Number of toys to learn shapes	73%	75%
Time investments	Number of different play activities	95%	98%
	Times told a story to child in last 3 days	67%	83%
	Times read to child in last 3 days	70%	85%
	Times played with child and toys in last 3 days	64%	81%
	Times labelled things to child in last 3 days	65%	82%
Mother's cognitive skills	Mothers' years of education*	64%	63%
	Mother's vocabulary	70%	69%
	Number of books for adults in the house*	40%	39%
	Number of magazines and newspapers	18%	17%
	Raven's score ("IQ") **	60%	59%
Mother's socio-emotional skills	Did you feel depressed? (-)	42%	46%
	Bothered by what usually don't bother you? (-)	28%	32%
	Had trouble keeping mind on doing? (-)	35%	38%
	Felt everything you did was an effort? (-)	31%	34%
	Did you feel fearful? (-)	24%	27%
	Did you sleep was restless? (-)	30%	34%
	Did you feel happy? (-)	13%	15%
	How often did you feel lonely last week? (-)	31%	35%
	Did you feel you couldn't get going? (-)	39%	42%

Note: This table shows the measures allowed to load on each latent factor, as well as the fraction of the variance in each measure that is explained by the variance in signal, for the control and treatment groups separately. Measures followed by an asterisks (*) were collected at baseline and measures followed by two asterisks (**) were collected at follow-up II. All other measures were collected at follow-up I at the end of the intervention. The sign (-) following measures of the child's non-cognitive skill and the mother's non-cognitive skill indicates that the scoring on these measures was reversed so that the corresponding latent factor is such that a higher score means a higher level of non-cognitive skill.

5.1 The measurement system and the distribution of factors

Table 2 describes the specification of the measurement system, which underlies the estimates of the production function, that is the set of variables used as measures of each latent factor. To arrive at this specification, we performed an exploratory factor analysis of the data that helped us to determine the number of factors that could be extracted from the data and to allocate measures to particular factors. The steps and results of this exploratory factor analysis are discussed in detail in Appendix C.

As mentioned above, identification requires that at least one measure for each factor is conditionally independent of the other measures for the same factor. In our case, this assumption can be justified by the fact that some developmental outcome variables are based on child level observations and are collected by a trained psychologist in community centers, while others are based on maternal reports and are collected in the home (on a different day) by an interviewer. The independence of measurement errors is probably not a far fetched assumption in such a context.

From the first stage of the estimation procedure, we obtained estimates of the measurement system, i.e. estimates of the mean and variance-covariance matrix of the latent factors for each group $d = \{0, 1\}$, estimates of the factor loadings and of the variances of the measurement error.[14] All these parameter estimates are reported in Appendix C.[15] Using these estimates, it is possible to

[14]We standardized all measures with respect to their mean in the control group, so we bypass the estimation of the intercepts.

[15]More precisely, Appendix Table A3 reports the estimates of the factor loadings in each measurement equation. Appendix Table A4 reports the estimates of the means of the latent factor distributions for the treated households relative to the means of the latent factor

assess the extent of measurement error in each measure of skill and investment in the data. Simplifying the notation, recall that the jth measure of latent factor θ can be written as:

$$m_j = \mu_j + \alpha_j \theta + \epsilon_j$$

Given this functional form and our maintained assumption that the factors are independent from the measurement error, the variance of m_j can be decomposed as: $Var(m_j) = \alpha_j^2 Var(\theta) + Var(\epsilon_j)$. We can then calculate the fractions of $Var(m_j)$ due to signal (s^θ) and noise (s^ϵ) as follows:

$$s_j^\theta = \frac{\alpha_j^2 Var(\theta)}{\alpha_j^2 Var(\theta) + Var(\epsilon_j)} \qquad \text{and} \qquad s_j^\epsilon = 1 - s_j^\theta \tag{9}$$

We report estimates of s^θ and s^ϵ in Table 2 for the control and treated groups. Clearly, there is much variation in the amount of information contained in each measure of the same factor. For example, 76% of the variance in Bayley Cognitive is due to signal, whereas only 38% of the variance in Mac Arthur-Bates Complex Sentences is due to signal. However, no measure is close to having 100% of its variance accounted for by signal, which emphasizes the importance of accounting for measurement error through the latent factor model.[16]

distributions for the control households (which is normalized to 0). Appendix Table A5 reports the estimates of the variance of and correlation between the latent factors for the treated and control households separately.

[16] There is little difference in the signal-to-noise ratio of measures between the control and treated groups. Under our assumption of invariant factor loadings and measurement error between the control and treated groups, differences in signal-to-noise ratios could only come from differences in the variances of the latent factors between the two groups. Appendix Table A5 shows that these are minimal.

Figure 1: Kernel densities of latent factors

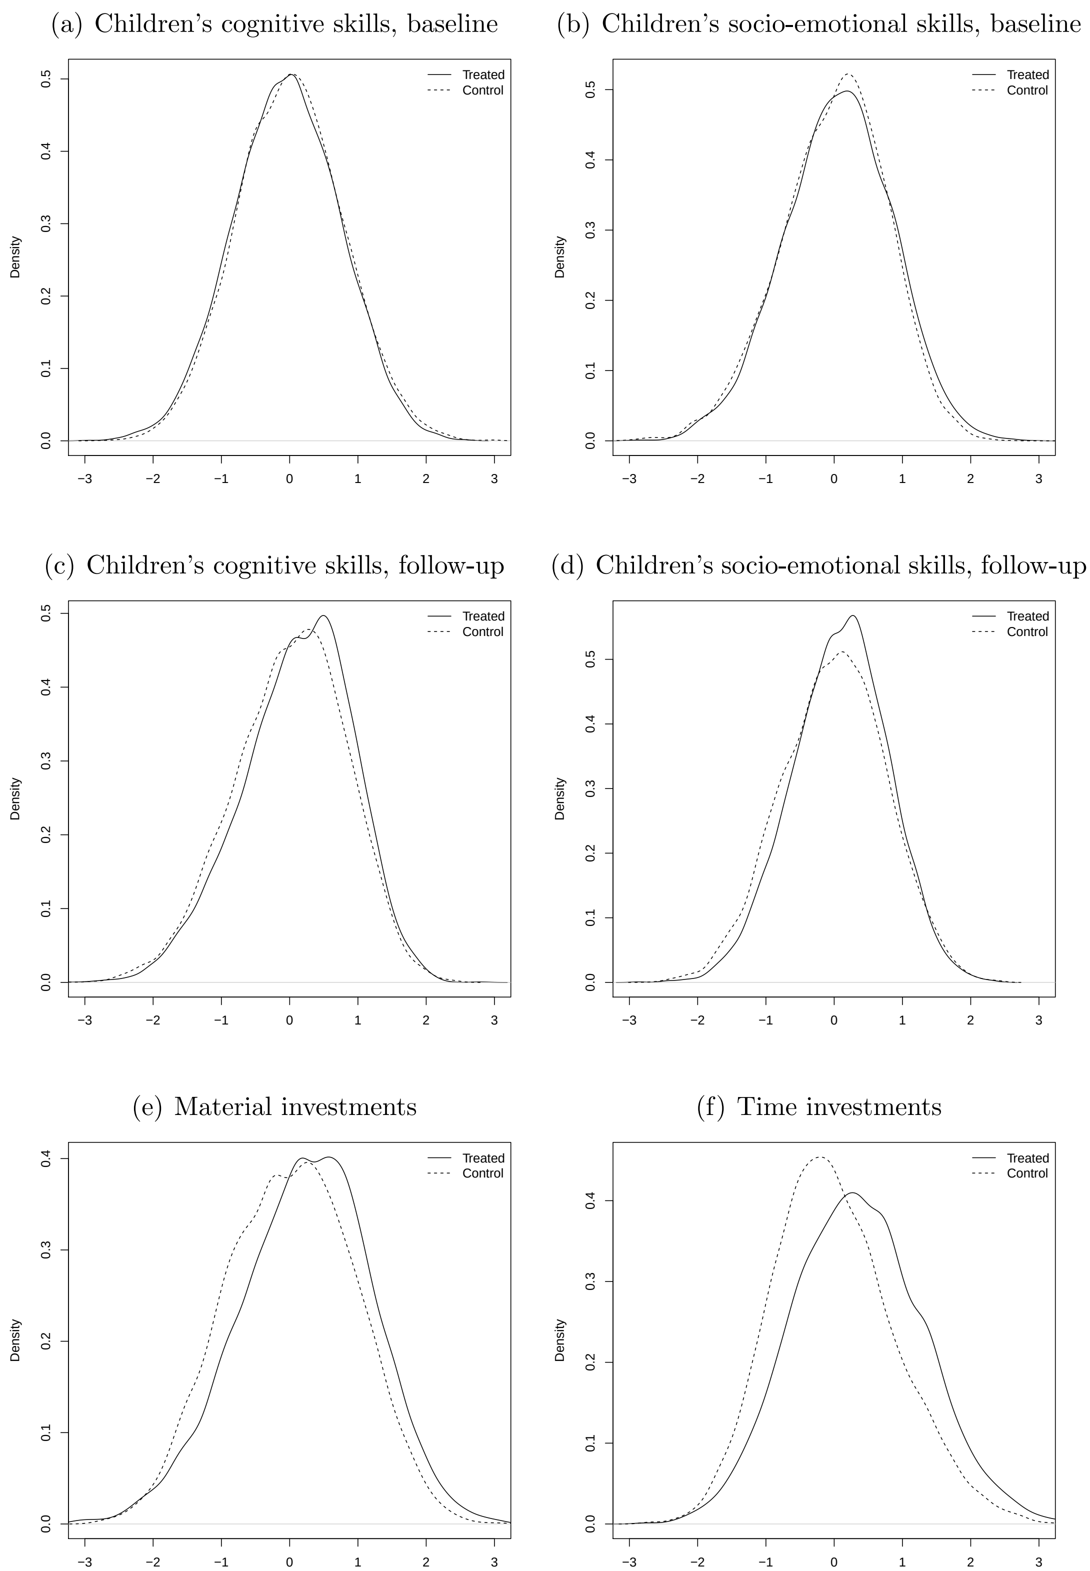

(a) Children's cognitive skills, baseline

(b) Children's socio-emotional skills, baseline

(c) Children's cognitive skills, follow-up

(d) Children's socio-emotional skills, follow-up

(e) Material investments

(f) Time investments

Note: These kernel densities are constructed using 10,000 draws from the estimated joint distribution of latent factors for the control group and for the treated group.

Having identified the entire distribution of factors for each group, we can study whether the intervention has changed the entire shape of these distributions, in addition to their means. In Figure 1, we plot the estimated Kernel densities of some of the factors. The first two panels show the distribution, in treatment and control villages, of cognitive and socio-emotional skills at baseline. These first two pictures confirm that our sample is substantially balanced. The following two panels depict the distribution of cognitive and socio-emotional factors at follow-up. In the case of cognitive factors we see that the shift in the mean reported in Appendix Table A4 reflects a shift in the entire distribution. For socio-emotional factors, however, the shift occurs mainly for children below the median.

Finally, in the last two panels, we notice a strong shift to the right of both the material and time investment factors. This suggests that at least part of the impact of the intervention is likely to have been driven by increases in both time and materials devoted by parents to the upbringing of their children.

5.2 Estimates of the investment functions

In Table 3, we present estimates of the investment equations. The first column presents the equation for material investments and the second column for time investments. As far as we know, our paper is unique in distinguishing between material and time investments in the context of estimating non-linear technologies of skill formation. Note that the results reported in Table 3 exclude interactions of the treatment parameter with the remaining variables. In earlier versions we found such interactions to be insignificant, i.e. the shift

Table 3: Estimates of the log-linear investment function

	Log of material investments	Log of time investments
Constant	**0.001**	**0.004**
	(0.016)	(0.015)
	[-0.025,0.027]	[-0.02,0.028]
Treatment dummy	**0.248**	**0.361**
	(0.073)	(0.065)
	[0.115,0.349]	[0.235,0.451]
Log of child's cognitive skills at *t*	**0.141**	**0.116**
	(0.061)	(0.057)
	[0.032,0.231]	[0.007,0.198]
Log of child's socio-emotional skills at *t*	**-0.008**	**0.031**
	(0.058)	(0.056)
	[-0.105,0.084]	[-0.053,0.13]
Log of mother's cognitive skills	**0.668**	**0.462**
	(0.082)	(0.079)
	[0.54,0.815]	[0.317,0.573]
Log of mother's socio-emotional skills	**-0.120**	**-0.310**
	(0.074)	(0.103)
	[-0.25,-0.001]	[-0.474,-0.123]
Log of wealth at *t*	**0.081**	**-0.086**
	(0.071)	(0.090)
	[-0.019,0.217]	[-0.231,0.06]
Mother is married at *t+1*	**0.126**	**0.115**
	(0.027)	(0.027)
	[0.075,0.164]	[0.066,0.155]
Log of number of children at *t+1*	**-0.096**	**-0.090**
	(0.033)	(0.033)
	[-0.146,-0.04]	[-0.151,-0.041]
Log of average male wages in village at *t+1*	**0.075**	**-0.026**
	(0.041)	(0.044)
	[-0.007,0.117]	[-0.106,0.039]
Log of average female wages in village at *t+1*	**0.004**	**0.033**
	(0.038)	(0.031)
	[-0.077,0.054]	[-0.013,0.088]

Note: Standard errors in parentheses and 90% confidence intervals in brackets are obtained using the non-parametric bootstrap described in Section 4. Appendix B provides a detailed description of the variables used to measure each latent factor.

in investment seems to have been uniform across groups with differing backgrounds.[17]

The first striking result is the impact of treatment on investments: it increases resources by 25% and time by 36% and both effects are highly significant. Thus, the intervention increased the time and the resources that parents provide to children. Referring back to the measurement system (Table 2), it is worthwhile noting that the time inputs are measured in a way that are targeted to child educational activities, such as the number of times an adult read to the child in the last three days. In other words, they do not refer simply to time spent with the child, but to interactions that promote development. Similarly, material investments refer to particular types of toys and play materials. Importantly, our estimates of the impact of the intervention on investments are uniquely driven by the experimental design and do not require any of the assumptions necessary for the identification of the production functions.

Turning now to the other regressors, we find that both time and material investments increase with the child's cognitive skills, but socio-emotional skills have no impact on investments, at least at the very young ages we are considering. The elasticity of both material and time investments with respect to maternal cognition is very high and particularly so for the former; however mother's socio-emotional skills have no significant effect. Married mothers in-

[17]Appendix Table A6 reports the estimates of the investment equation in which we interact all the regressors with the treatment dummy. We test the joint significance of the interaction terms and find that we cannot reject that all the interactions are equal to 0 for both material and time investments at conventional levels of significance (the p-value for the material investment equation is 0.577 and the p-value for the time investment equation is 0.667). Appendix Table A7 reports the estimates of the production functions when the control functions are derived from the fully interacted investment function. We observe no major difference between these results and the ones based on the non-interacted investment function reported in Tables 4 and 5.

vest more time and more materials, but the overall number of children reduces investments. We find no effect of wealth on investments, probably because our sample is quite homogeneously poor. However, we do find that more is invested in children in terms of materials when male wages are higher in the village, which we interpret as an income effect. Average female wages in the village, which could reflect the opportunity cost of time for the mother as well as her access to resources, have no effect.

Of the above variables, wealth, marital status, and male and female wages at the village level are excluded from the production function and serve as instruments. We test for the joint significance of the exclusion restrictions in the material and time investment equations. The p-value from a F-test performed separately on each equation is 0.008 for the material investment equation and 0.041 for the time investment equation. The p-value of a test that the rank of the coefficient matrix of the excluded variables across both investment equations is less than 2 has a p-value of 0.035.[18] The fact that some of the coefficients on the individual variables are not statistically different from zero is probably driven by collinearity and by the limited variation we have (as some of these variables only vary across village). From the high joint significance of our instruments however, we conclude that they are strong enough to identify the impact of investment in the production function for skills, to which we now turn.

[18] To perform this rank test, we proceeded as follows. Let B be the 2×4 matrix of coefficients on the four excluded variables in the two investment equations. We used the non-parametric bootstrap to derive critical values for a test that the smallest eigenvalue of $B'B$ is zero. See Robin and Smith (2000) for a discussion of such tests and Blundell, Duncan, and Meghir (1998) for a similar application.

5.3 Estimates of the production functions

Tables 4 and 5 report the estimates of the CES production function. For each skill, we report estimates of the parameters of the production function obtained both ignoring the endogeneity of investments (in the left column) and taking it into account by including the residuals from the investment functions as additional inputs (in the right column), as specified in equation (4). It should be stressed that very limited systematic evidence exists on the evolution and development of children skills at such young ages.

We have experimented with allowing all coefficients to vary with treatment status, but we found no evidence that the production function differs systematically as a result of the intervention. In fact the tests that the coefficients are the same across intervention and control groups, other than total factor productivity A, have p-values of 0.879 and 0.889 for the production functions of cognitive and socio-emotional skills, respectively.[19] Indeed, as Tables 4 and 5 show, even total factor productivity remains unaffected by treatment, except when we treat investments as exogenous in the production function for cognitive skills (first column of Table 4).[20]

The results provide evidence of several important features of skill development. First, we find strong evidence of self-productivity of skills. That is, the current stock of cognitive (socio-emotional) skills strongly affects the

[19] Appendix Tables A8 and A9 report the estimates of the production function for each skill type in which we allow all parameters (except the elasticity of substitution) to vary between treated and controls. None of the interactions, except that with the number of children in the household, is significant.

[20] We also experimented with a specification of the model where we allowed the child's age at follow-up to enter the investment equations and to interact with the total-factor productivity parameter (A) in the production functions. We did not find evidence that this parameter depended on the child's age however (see Appendix Tables A10 and A11.)

Table 4: Estimates of the CES production function for cognitive skills

	Without control function	*With control function*
Child's cognitive skills at *t*	**0.591**	**0.566**
	(0.043)	(0.057)
	[0.527,0.67]	[0.489,0.674]
Child's socio-emotional skills at *t*	**0.03**	**0.038**
	(0.043)	(0.050)
	[-0.037,0.106]	[-0.035,0.126]
Mother's cognitive skills	**0.194**	**0.037**
	(0.049)	(0.131)
	[0.107,0.264]	[-0.194,0.223]
Mother's socio-emotional skills	**0.06**	**0.051**
	(0.045)	(0.049)
	[-0.016,0.126]	[-0.028,0.127]
Material investments at *t+1*	**0.082**	**0.397**
	(0.033)	(0.208)
	[0.036,0.144]	[0.128,0.765]
Time investments at *t+1*	**0.008**	**-0.138**
	(0.035)	(0.142)
	[-0.056,0.057]	[-0.421,0.039]
Number of children in household at *t+1*	**0.035**	**0.049**
	(0.026)	(0.030)
	[-0.009,0.076]	[0.002,0.1]
Control function for material investments	**-**	**-0.33**
		(0.218)
		[-0.715,-0.023]
Control function for time investments	**-**	**0.156**
		(0.151)
		[-0.037,0.453]
Complementarity parameter	**0.123**	**0.07**
	(0.082)	(0.060)
	[-0.025,0.243]	[-0.032,0.161]
Elasticity of substitution	**1.141**	**1.075**
	(0.106)	(0.070)
	[0.976,1.321]	[0.969,1.192]
Productivity parameter (A)	**0.984**	**0.993**
	(0.012)	(0.011)
	[0.966,1.005]	[0.972,1.008]
Productivity parameter interacted with treatment	**0.1**	**0.08**
	(0.052)	(0.072)
	[0.028,0.198]	[-0.012,0.228]

Note: Standard errors in parentheses and 90% confidence intervals in brackets are obtained using the non-parametric bootstrap described in Section 4. Appendix B provides a detailed description of the variables used to measure each latent factor.

Table 5: Estimates of the CES production function for socio-emotional skills

	Without control function	*With control function*
Child's cognitive skills at *t*	**0.11**	**0.122**
	(0.044)	(0.059)
	[0.039,0.185]	[0.024,0.222]
Child's socio-emotional skills at *t*	**0.435**	**0.413**
	(0.055)	(0.059)
	[0.374,0.552]	[0.354,0.537]
Mother's cognitive skills	**-0.054**	**0.116**
	(0.066)	(0.142)
	[-0.168,0.046]	[-0.201,0.276]
Mother's socio-emotional skills	**0.151**	**0.161**
	(0.058)	(0.058)
	[0.047,0.233]	[0.046,0.235]
Material investments at *t+1*	**0.14**	**-0.32**
	(0.043)	(0.198)
	[0.079,0.219]	[-0.529,0.108]
Time investments at *t+1*	**0.119**	**0.434**
	(0.041)	(0.133)
	[0.043,0.181]	[0.17,0.591]
Number of children in household at *t+1*	**0.099**	**0.073**
	(0.026)	(0.027)
	[0.048,0.136]	[0.025,0.113]
Control function for material investments	**-**	**0.477**
		(0.204)
		[0.043,0.711]
Control function for time investments	**-**	**-0.336**
		(0.136)
		[-0.506,-0.068]
Complementarity parameter	**0.049**	**0.006**
	(0.077)	(0.056)
	[-0.085,0.158]	[-0.059,0.12]
Elasticity of substitution	**1.051**	**1.006**
	(0.088)	(0.063)
	[0.921,1.187]	[0.944,1.137]
Productivity parameter (A)	**0.987**	**0.992**
	(0.016)	(0.012)
	[0.966,1.019]	[0.976,1.014]
Productivity parameter interacted with treatment	**-0.015**	**-0.016**
	(0.042)	(0.058)
	[-0.081,0.054]	[-0.115,0.073]

Note: Standard errors in parentheses and 90% confidence intervals in brackets are obtained using the non-parametric bootstrap described in Section 4. Appendix B provides a detailed descriptions of the variables used to measure each of the latent factor.

development of future (socio-emotional) cognitive skills. Second, we find evidence of cross-productivity in the production function for socio-emotional skills (Table 5): the current stock of cognitive skills fosters the development of future socio-emotional skills. Although the current stock of socio-emotional skills has a positive coefficient in the production function of cognitive skills, the estimate is not significantly different from zero and we cannot conclude there is cross-productivity in the production of cognitive skills. This result contrasts with that reported by Cunha, Heckman, and Schennach (2010), who find socio-emotional skills to be important for the accumulation of future cognitive skills. Our evidence, however, relates to a different, much earlier, period of the life-cycle.

Our results also provide evidence that parental investments play a crucial role in the development of both cognitive and socio-emotional skills and that they are endogenous. With respect to cognitive skills, the results reveal that only material investments seem to matter; their effect becomes very large and strongly significant, once we control for endogenous investments. With respect to socio-emotional skills, when we treat investments as endogenous, only time investments significantly matter for the accumulation of socio-emotional skills. The point estimate of the share parameter for material investments is negative, but very imprecise. The effect of time investments on future socio-emotional skills is very large and strongly significant.

Investments are also clearly endogenous. The control function in the cognitive production function is highly significant for material investments and including it increases the share parameter for material investments more than fivefold. Similarly the control function for time investments in the corresponding production function is also highly significant, and including it in the production function triples the coefficient on time investments. In both of these cases, the estimate of the coefficient of the relevant control function is negative, suggesting that parents compensate with materials and time when adverse events affect the development of their child. This evidence is important as it speaks to the issue of whether parental investment reinforces or compensates shocks received by children, an issue discussed by a large literature, some of which is surveyed in Almond and Mazumder (2013). Our findings indicate very clearly that investment seems to be compensating rather than reinforcing.

The estimates of the production functions show that mothers' cognitive skills do not seem to have a significant *direct* effect on either children's cognitive or socio-emotional development, once we control for the endogeneity of investments. However, it should be remembered that according to our estimates of the investment function (Table 3), mother's cognitive skills have an important effect on both types of investment and, therefore, an indirect effect on child development. Mother's socio-emotional skills have a significant impact on the development of the child's socio-emotional skills, even after controlling for the endogeneity of investments, but their effect on the child's future cognitive skills is not significantly different from zero.

It is also interesting that the number of siblings improves outcomes for children *conditional* on investments, presumably because of mutual stimulation and interaction. On the other hand, the estimates of the investment function imply that siblings also dilute investments to the subject child, thus creating two counteracting effects.

The last key component of the production function is the complementarity parameter, ρ_k. In both production functions, it is estimated to be zero, which implies an elasticity of substitution equal to one. That is, the production function for cognitive and socio-emotional skills has a Cobb-Douglas form.

5.4 Implications for human capital accumulation and the role of the intervention

The results presented above have interesting implications for human capital accumulation. First, the production function is Cobb-Douglas and hence inputs are complementary. Moreover, baseline skills, investments and, to some extent, maternal skills are important for the development of future skills. Together, these two facts imply that investments are more productive for children with higher levels of early cognition and socio-emotional skills.

We illustrate this point in Figure 2, which plots the marginal product of investments as a function of the child's baseline skills. Consider first Figure 2a. The productivity of material investment in the production function for cognitive skills triples as we move from the lower deciles to the upper deciles of the distribution of baseline cognition. Similarly, time investments are more productive for the accumulation of socio-emotional skills, the higher a child is on the distribution of socio-emotional skills at baseline (Figure 2b). This complementarity illustrates both the importance of targeting low-achieving children early as well as the difficulty of delivering interventions with high returns for them, particularly as time passes.

Figure 2: Complementarity between investments and baseline skills

(a) Marginal product of material investments in the production of cognitive skills

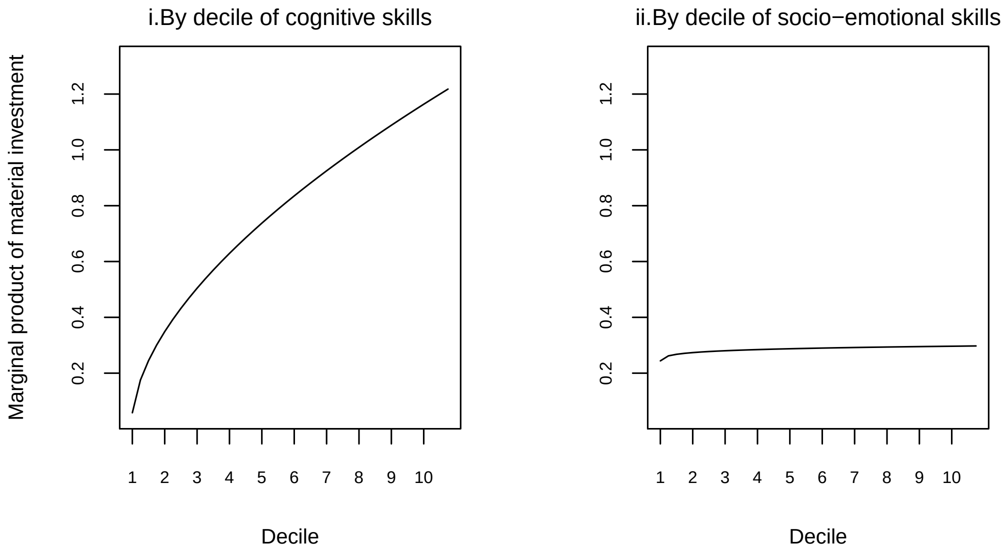

(b) Marginal product of time investments in the production of socio-emotional skills

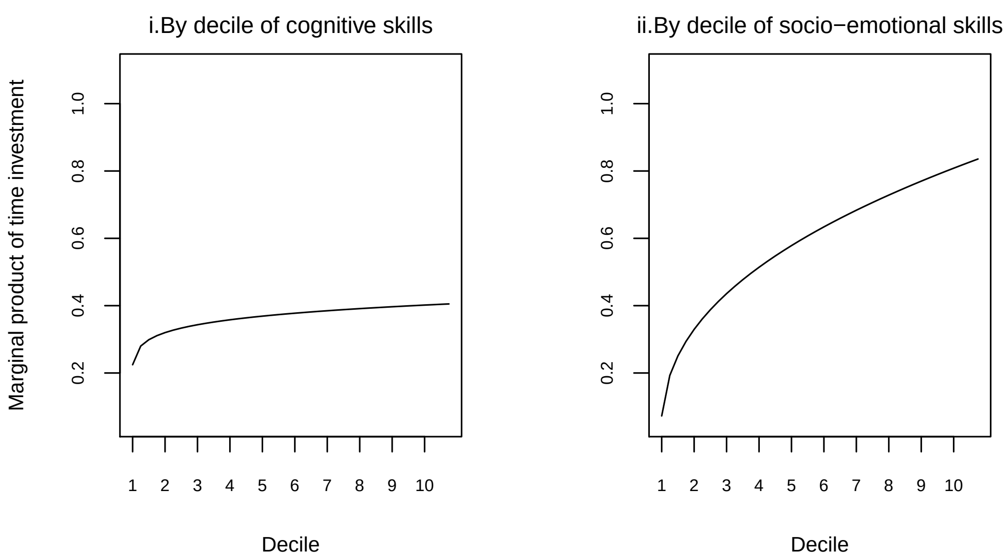

Notes: The marginal product of material and time investments in the production function for skill of type k are given by: $A_d^k \gamma_{5,d}^k I_{i,t}^{M\,(\rho_k-1)} Q^{(\frac{1-\rho_k}{\rho_k})}$ and $A_d^k \gamma_{6,d}^k I_{i,t}^{T\,(\rho_k-1)} Q^{(\frac{1-\rho_k}{\rho_k})}$, respectively, where $Q = [\gamma_{1,d}^k(\theta_{i,t}^C)^{\rho_k} + \gamma_{2,d}^k(\theta_{i,t}^S)^{\rho_k} + \gamma_{3,d}^k(P_i^C)^{\rho_k} + \gamma_{4,d}^k(P_i^S)^{\rho_k} + \gamma_{5,d}^k(I_{i,t}^M)^{\rho_k} + \gamma_{6,d}^K(I_{i,t}^T)^{\rho_k}]$. The figures above are constructed by evaluating these objects at different deciles of $\theta_{i,t}^C$ (figures i) and $\theta_{i,t}^S$ (figures ii) and holding all remaining inputs of the production function at their mean values pooling observations from the control and treatment groups together.

As we mentioned above, the factor-neutral productivity parameter A is not shifted by the treatment indicator, whose coefficient is small and insignificant in either production function. Moreover, when we test whether the intervention has shifted any of the other parameters of the production function, we cannot reject the null of no change. On the other hand, there is a positive and significant shift in parental investments among treated families relative to non-treated families (Figure 1), and on average there is a 25% increase in material investments and 36% increase in time investments among treated households relative to control households. These results imply that the intervention acted exclusively by changing parental investments in materials and time and not by increasing the productivity of investments nor that of any other inputs. Given the evidence we just presented on the role that these investments play in the production function for human capital, it is natural to ask whether this channel can explain all the impact of the intervention.

More formally, we ask what fraction of the difference in skills between the treated and non-treated children can be explained by the change in parental investments induced by the intervention. We can answer this question by using the structure and estimates of the model. In particular, we predict the distributions of future cognitive and socio-emotional skills given the distribution of investments in the treatment group and those of the control group respectively. We then compute the gap between the two for each type of skill. This difference measures how much of the impact of the intervention we can explain using just the change in investments.

Table 6 reports the results from this exercise based on the estimates with

Table 6: Decomposition of the treatment effects in skills

	Gap in log skills between treated and controls	
	Cognitive skills at *t+1*	Socio-emotional skills at *t+1*
Data	0.115	0.075
	(0.056)	(0.045)
	[0.011,0.195]	[0.007,0.151]
Model	0.107	0.088
	(0.058)	(0.047)
	[0.009,0.195]	[0.005,0.156]

Note: Standard errors in parentheses and 90% confidence intervals in brackets are obtained using the non-parametric bootstrap described in Section 4. Appendix B provides a detailed description of the variables used to measure each latent factor.

endogenous investments. The results indicate that the model predicts the actual gap in log skills between the treated and non-treated children well for both type of skills. Specifically the model predicts 93% of the cognitive improvement and slightly overshoots the socio-emotional improvement among treated children. The differences between the data and the model predictions are not significant. Thus the conclusion is quite clear: the intervention changed parental behavior by inducing more of the appropriate investments but did not change the production function.

An interpretation of this finding is that, as a result of the intervention, parents' behavior did not become more productive in itself, but parents allocated more of their resources on their children, given their beliefs on the effectiveness of these investments. This suggests that a reason for under-investment in child development among poor populations might be a lack of knowledge about the effectiveness of investments, thus implying that one way interventions can be successful is by changing parenting beliefs on the subject.

6 Conclusion

In this paper, we estimate the production functions for cognitive and socio-emotional skills in early childhood and analyze the mechanisms through which a randomized early childhood intervention in Colombia affected child development. The intervention itself - a psychosocial stimulation program delivered through weekly home visits - significantly improved the cognitive and socio-emotional development of children aged between 12 and 24 months at the start of the intervention. Underlying our approach is a model of parental investments in children, constrained by a production function for cognitive and socio-emotional skills.

We make three interrelated contributions. First, we estimate equations for material and time investments depending on both family and child baseline characteristics. Second, we estimate production functions for cognitive and socio-emotional skills that depend on the investments, which are taken to be endogenous. Third, we use this framework to trace out the way the intervention affected the production of children's skills.

In line with the existing literature, we find strong evidence that a child's current stock of skills fosters the development of future skills (self-productivity). We also find evidence that early cognition helps improve socio-emotional skills (cross-productivity). Second, and also in line with the existing literature, we find that current skills, parental investments and maternal human capital are complementary in the production of future skills. This means, for example, that parental investments are more productive when the child's current skills are higher. Lastly, we find that parental investments greatly matter for the accumulation of cognitive and socio-emotional skills. In particular, material investments seem to matter more for cognitive skills, while time investments seem to matter more for socio-emotional skills. Our paper is the first to estimate a non-linear production function for skills with different types of investments, and our results show that it is a non-trivial extension of the standard model.

When we compare the production function estimates between treatment and control groups, we cannot reject that they are the same. We find that we can attribute the entire effect of the intervention to an increase in time and material resources devoted by parents to their children. In other words, the intervention did not shift the way that skills are produced. Nevertheless, we show that even a shift in investments without a change in the production function succeeded in altering children's outcomes in significant ways. Whether the intervention changed parental behavior fundamentally enough to yield long-lasting impacts is an open question, but one of crucial importance as further follow-up data is collected. The framework developed here surely offers a valuable tool to understand what drives the sustainability of this and other early years interventions and organize the evidence in a systematic way.

References

Almond, D. and Currie, J. (2011). Human capital development before age five. In O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics, Volume 4b*. Elsevier Science, Amsterdam.

Almond, D. and Mazumder, B. (2013). Fetal origins and parental responses. *The Annual Review of Economics* 5(3):1–20.

Anderson, T. and Rubin, H. (1956). Statistical inference in factor analysis. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, volume 5, 111–150. University of California Press, Berkeley.

Attanasio, O., Fernandez, C., Fitzsimons, E., Grantham-McGregor, S., Meghir, C., and Rubio-Codina, M. (2014). Using the infrastucture of a Conditional Cash Transfer programme to deliver a scalable integrated early child development programme in Colombia: A cluster randomised controlled trial. Forthcoming in *British Medical Journal* 349:g5785.

Attanasio, O., Meghir, C., and Nix, E. (2015). Investments in children and the development of cognition and health in India. Mimeo, Yale University.

Attanasio, O., Meghir, C., and Santiago, A. (2012). Education choices in Mexico: Using a structural model and a randomized experiment to evaluate Progresa. *Review of Economic Studies* 79(1):37–66.

Bates, J., Freeland, C., and Lounsbury, M. (1979). Measurement of infant difficultness. *Child Development* 50:794–803.

Bayley, N. (2006). *Bayley Scales of Infant and Toddler Development*. Harcourt Assessment, San Antonio, TX, 3rd edition.

Blundell, R., Duncan, A., and Meghir, C. (1998). Estimating labor supply responses using tax reforms. *Econometrica* 66(4):827–761.

Carneiro, P., Hansen, K., and Heckman, J. (2003). Estimating distribution of counterfactuals with an application to the returns to schooling and measurement of the effect of uncertainty on schooling choice. *International Economic Review* 44(2):361–422.

Cunha, F. and Heckman, J. (2008). Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of Human Resources* 43(4):738–782.

Cunha, F., Heckman, J., Lochner, L., and Masterov, D. (2006). Interpreting the evidence on life cycle skill formation. In E. Hanushek, S. Machin, and L. Woessmann, eds., *Handbook of the Economics of Education, Volume 1*. Elsevier Science, Amsterdam.

Cunha, F., Heckman, J., and Schennach, S. (2010). Estimating the technology of cognitive and non-cognitive skill formation. *Econometrica* 78(3):883–931.

Duflo, E., Hanna, R., and Ryan, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review* 102(4):1241–78.

Frongillo, E., Sywulka, S., and Kariger, P. (2003). UNICEF psychosocial care indicators project. Final report to UNICEF. Mimeo, Cornell University.

Gertler, P., Heckman, J., Pinto, R., Zanolini, A., Vermeerch, C., Walker, S., Chang, S., and Grantham-McGregor, S. (2014). Labor market returns to an early childhood stimulation intervention in Jamaica. Forthcoming, *Science*.

Gorusch, R. (1983). *Factor Analysis*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Gorusch, R. (2003). Factor analysis. In *Handbook of psychology: Research methods in psychology*, volume 2, 143–164. John Wiely & Sons, Inc., Hoboken, NJ.

Grantham-McGregor, S., Powell, C., Walker, S., and Himes, J. (1991). Nutritional supplementation, psychosocial stimulation, and mental development of stunted children: the Jamaican study. *Lancet* 338(758):1–5.

Heckman, J., Moon, S., Pinto, R., Savelyev, P., and Yavitz, A. (2010). Analyzing social experiments as implemented: A reexamination of the evidence from the HighScope Perry Preschool Program. *Journal of Quantitative Economics* 1:1–46.

Heckman, J., Pinto, R., and Savelyev, P. (2013). Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *American Economic Review* 103:2052–2086.

Helmers, C. and Patnam, M. (2011). The formation and evolution of childhood skill acquisition: Evidence from India. *Journal of Development Economics* 95(2):252–266.

Jackson-Maldonado, D., Marchman, V., and Fernald, L. (2012). Short form versions of the Spanish MacArthur-Bates Communicative Development Inventories. In *Applied Psycholinguistics*. Cambridge University Press, Cambridge.

Knudsen, E. (2004). Sensitive periods in the development of the brain and behavior. *Journal of Cognitive Neuroscience* 16(1):1412–1425.

Knudsen, E., Heckman, J., Cameron, J., and Shonkoff, J. (2006). Economic, neurobiological and behavioral perspectives on building America's future workforce. *Proceedings of the National Academy of Science* 103(27):10155–10162.

McGregor, S. G., Cheung, Y., Cueto, S., Glewwe, P., Richter, L., and Strupp, B. (2007). Developmental potential in the first 5 years for children in developing countries. *The Lancet* 369(9555):60–70.

Putnam, S. P., Gartstein, M., and Rothbart, M. K. (2006). Measurement of fine-grained aspects of toddler temperament: The Early Childhood Behavior Questionnaire. *Infant Behavior and Development* 29(3):386–401.

Robin, J. and Smith, R. (2000). Tests of rank. *Econometric Theory* 16:151–175.

Todd, P. and Wolpin, K. (2006). Using experimental data to validate a dynamic behavioral model of child schooling and fertility: Assessing the impact of a school subsidy program in Mexico. *American Economic Review* 96(5):1384–1417.

Walker, S., Chang, S., C.Powell, E.Simonoff, and Grantham-McGregor, S. (2006). Effects of psychosocial stimulation and dietary supplementation in early childhood on psychosocial functioning in late adolescence: follow-up of randomized controlled trial. *British Medical Journal* 333(7566):472.

Walker, S., Chang, S., Vera-Hernandez, M., and Grantham-McGregor, S. (2011). Early childhood stimulation benefits adult competence and reduces violent behavior. *Pediatrics* 127(5):849–857.