

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, Revised May 2017

We are grateful to three anonymous referees and the editor for detailed comments. We also 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. We thank the NIH for funding under grant R01 HD072120 as well as 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>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Orazio Attanasio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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, Revised May 2017

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 aged 12 to 24 months at baseline. We estimate the determinants of material and time investments in these children and evaluate the impact of the treatment on such investments. We then estimate the production functions for cognitive and socio-emotional skills. The effects of the program can be explained by increases in parental investments, which have strong effects on outcomes and are complementary to both maternal skills and child's baseline skills.

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

Sarah Cattan
Institute for Fiscal Studies
7 Ridgmount Street London
WC1E 7AE
United Kingdom
sarah_c@ifs.org.uk

Marta Rubio-Codina
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
United Kingdom
marta_r@ifs.org.uk

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

A 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 play a crucial role in the process of human capital accumulation. At this time however, many children are exposed to risk factors such as poverty, malnutrition and non-stimulating home environments preventing them from reaching their full potential, particularly so in developing countries (Grantham-McGregor et al. (2007), Lu et al. (2016) and Black et al. (2016)). These factors are likely to play an important role in the intergenerational transmission of poverty.

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 (Grantham-McGregor et al. (1991), Walker et al. (2011) and Gertler et al. (2014)), the Perry Preschool program (Heckman et al., 2010) and the Abecedarian experiment (Campbell and Ramey (1994), Campbell et al. (2014)). In Attanasio et al. (2014), we present the impacts of an 18-month long early childhood intervention in Colombia targeted at disadvantaged children aged 12-24 months old at baseline. The intervention was based on the Jamaican model in that it offered psycho-social stimulation via weekly home visits and micronutrient supplementation. However, it was designed to be scalable by training local women involved in the administration of a large welfare program to administer the weekly home visits. The randomized controlled trial (RCT) used to evaluate each arm of the intervention showed that stimulation led to highly significant improvements in cognition and language development measured immediately following the end of the intervention.¹ Micronutrient supplementation did not affect any outcome observed in the data.

Building on these results, the main aim of this paper is to understand *how* the stimulation

¹Cognition improved by 26% of a Standard Deviation (SD) (p-value 0.002) and receptive language by 22% of a SD (p-value 0.032).

component of the intervention led to improvements in child development. The intervention could have affected child development through different channels. For example, it could have led parents to make greater material and time investments in their children. But it could also have changed the production function for child skills, through the direct effect of the home visits as a new input or by changing the effectiveness of parental inputs. In what follows, we build a model of parental investments and child skill formation to tease out the relative importance of these different mechanisms, a crucial step to better focus and increase the sustainability of interventions in the future.

We start by estimating the determinants of parental investments and assessing how the intervention changed parental choices. Indeed the way parents respond to such programs, which can be seen as a type of in-kind transfer, is an open question: the intervention could lead parents to reinforce their engagement with the child or instead crowd-out their investments. Gelber and Isen (2010), for example, provide evidence that the US early childhood program Head Start led to an increase in parental involvement, thus crowding-in household resources. In our treatment of the question here, we exploit the experimental variation induced by the RCT and distinguish between material investments (i.e. books and toys around the house) and time investments (i.e. time spent by an adult in the household on education activities with the child).² We show that parents increased both types of investment substantially as a result of the home visits.

We then estimate production functions for child cognitive and socio-emotional skills. The main inputs we specify are baseline child skills, maternal skills, and material and time investments. This technology is non-linear and allows the degree of substitutability between inputs to be determined from the data. Within this framework, we quantify how much changes in parental investments contributed to improving child outcomes in the treatment group. We also test whether the intervention changed the parameters of the production

²See DelBoca et al. (2014) for a structural model of household choices and child development based on the PSID Child Development Supplement data and also including time and resource investments.

function, which, as discussed above, could reflect the direct effect of the stimulation provided by the home visitors or a change in the productivity of inputs.

The two waves of data we use were collected just before and just after the intervention and contain rich measures of child development, maternal skills and parental investments. Importantly, we collect information on materials and activities that have an educational aspect, thus enabling a clear interpretation of parental behaviour as investments in their children. 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 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). We explore the sensitivity of the results to the possible econometric endogeneity of investments by implementing a control function approach as in Attanasio, Meghir, and Nix (2015), whose estimation procedure we adopt here.

The estimates of the investment functions reveal important information about some of the drivers of developmental inequality: children with better initial cognitive skills receive more investments and, crucially, mothers with higher skill levels invest more in their children *given* the child's skills. Our estimates of the production functions also reveal a series of interesting and important patterns. 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. Cognitive and socio-emotional skills are both important for developing future socio-emotional skills.³ Second, parental investments matter for the accumulation of skills. In particular, material investments seem to matter more for cognitive skills, while both material and time investments are important for socio-emotional skills. When we allow for investments to be endogenous, we estimate their impacts to be larger and this is in line with results obtained

³These features of the technology of skill formation are often referred to as *self-productivity* and *cross-productivity* (Cunha et al., 2006).

by Cunha, Heckman, and Schennach (2010) and Attanasio, Meghir, and Nix (2015), yet in very different contexts.⁴

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, the intervention does not seem to have changed any of the parameters of the production function. These two findings mean that the gains in cognitive and socio-emotional skills among children who received the intervention are mainly explained by changes in parental investments and imply that having the home visitor merely interact with the child for an hour a week, without trying to strengthen parenting practices, would have been unlikely to benefit children. This emphasizes the key importance of the parenting component of the intervention.

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. While there is a large literature evaluating the impact of early childhood interventions on child development, our paper innovates by complementing the information obtained from the RCT of a specific intervention with a 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),

⁴The 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). See also Helmers and Patnam (2011) for the estimation of a linear production function in India. Finally, and also in line with the existing literature, we find that current skills and parental investments are complementary in the production of future skills, meaning that returns to investments are higher for children with better initial conditions.

who document 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. Instead, we use a model in which parents make investment choices and human capital accumulates according to a production function, so as to interpret and explain the impacts induced by a successful intervention. Moreover, unlike the Jamaican intervention, which targeted malnourished children and the Perry Pre-School Program, which targeted children with specifically low cognition, we target a broader population. Our subjects are drawn from the beneficiaries of the Colombian conditional cash transfer (CCT) program *Familias en Acción*, which covers the poorest 20% of the population.⁵ In this sense, our program has the potential to serve as a model for early childhood policy that could be broadly implemented alongside CCT programs or other welfare programs targeting poor families.

The paper proceeds as follows. Section 2 provides some background on the intervention. Section 3 describes the data and the factor model approach we take to reduce their dimensionality and extract error-free measures of children’s skills, parental skills and investments. Section 4 discusses the short-term impacts of the intervention and some suggestive evidence of the underlying mechanisms. Section 5 presents the theoretical framework we use and discusses its empirical implementation. Section 6 presents the estimates of the model and discusses their implications for our understanding of the intervention. Section 7 concludes.

2 Background on the intervention and its evaluation

The 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

⁵See Attanasio et al. (2010) for a description and evaluation of that program.

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 around the size of California. It had two components: psychosocial stimulation and micronutrient supplementation. The stimulation curriculum was based on the Jamaican home visiting model, which obtained positive short- and long-term effects (Grantham-McGregor et al. (1991), Walker et al. (2006, 2011) and Gertler et al. (2014)). 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 as learning tools.

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, covering a large part of the country and obtaining much larger sample sizes. The second was that the intervention was designed to be scalable. To this end, home visitors were drawn from a network of local women, generated by the administrative set-up of the CCT program. Every 50-60 beneficiaries elect a representative who is in charge of organizing social activities and acts as mediators between them and the program administrators. These women, known as *Madre Líderes* (MLs), are beneficiaries of the program themselves and given they are selected by their peers one can deduce that they tend to be more entrepreneurial and proactive than the average beneficiary. In terms of specific characteristics they are about 10 years older and have about one more year of education than the subject mothers. Finally, as mentioned in the introduction, another distinct feature of our intervention that we targeted a more general poor population, namely the beneficiaries of the Colombian CCT program, as compared to the extreme disadvantage of the malnourished population targeted by the Jamaican experiment.

The intervention was evaluated through a cluster randomized controlled trial involving the random allocation of 96 municipalities. After first stratifying into three large regions, 32 municipalities in each were randomly assigned to one of 4 groups: (i) psychosocial stimulation, (ii) micronutrient supplementation, (iii) both, and (iv) control. In each municipality, 3 MLs were selected and the children aged 12-24 months of the beneficiary households represented by each of these MLs were recruited to the study. There was a total of 1,429 children living in 96 towns in central Colombia. Possibly because the ML are such trusted figures in their communities, compliance was high and the average number of home visits made was 63, which is 81% of those scheduled. The attrition rate between baseline and follow-up was around 10% across treatment arms, and the difference in loss among the groups was not statistically significant.⁶

As reported in Attanasio et al. (2014), there were no significant impact of micro-nutrient supplementation on any child developmental outcomes so, in this paper, we focus on the psychosocial stimulation arm of the program. In what follows therefore, we refer to the “treated” group as those children who received the stimulation component of the intervention (groups i and iii) and to the “control” group as those children who did not (groups ii and iv).

Individuals randomized into our intervention were all eligible for and receiving subsidies from the CCT program, which covers the 20% poorest in Colombia. On average, households were part of the program for 21 months at baseline. This feature is common between treatment and control communities, but it is true that the context in which our program was implemented and in particular the existence of the CCT may be a factor in how effective the program was. This, of course, is related to the more general issue of extrapolating the effects of the program to other contexts outside the support of the data. Nevertheless, CCT programs are quite common in low-and-middle income countries and consequently the context

⁶As we explain in Section 3.1, our data at baseline and at follow-up come from a household survey and from direct assessments administered to children in a community centre. The attrition rate for the household survey was 6.9%. The attrition rate for the direct assessments was 10.7%.

is directly relevant to many other countries besides Colombia.

Finally, a frequently asked question is whether the intervention is just “teaching to the test” without leading to genuine advances in cognition. First, implementation of the curriculum has been shown to have long-run effects on cognition (Walker et al., 2005, 2011) and labour market outcomes (Gertler et al., 2014). This in itself is evidence that it can induce deep changes in achievement rather than just teach children to remember a few activities and perform better on a test. More generally, the intervention curriculum emphasizes cognitive, language and socio-emotional development through play and the promotion of mother-child interactions. While some of the play activities specifically address the type of cognitive and fine motor skills (building towers with blocks, tracing lines) and concepts (shapes, sizes, colors) that are assessed in developmental tests, the focus is on learning through play in a supportive and stimulating environment. Activities are introduced progressively and in developmental order to facilitate scaffolding - i.e. increasing or decreasing the challenge based on the child’s performance - and there is a strong emphasis on praising attempts and not only successes. This approach and focus are aimed at promoting attention to task (or attention focusing), perseverance and self-esteem, which are also important skills required to perform well in developmental assessments (beyond the mastering of the concept or the task). Similarly, there is also a strong focus on labelling the environment and looking at picture books together, which are activities that enrich vocabulary and promote bonding, attention (i.e. following a story) and other cognitive abilities (linking concepts, understanding cause and effect relations). All of these skills are associated with improved school readiness, school attainment and other outcomes associated with socio-economic success in life.

3 Data and measurement system

3.1 Data

The data we use in this paper comes from two rounds of data collection: before the intervention started (baseline) and just after it ended 18 months later (follow-up). In each round, information was collected in two ways: via a household survey in the home and via tests directly administered to children in a community centre.

The household surveys contain information on an extensive set of socio-economic and demographic characteristics, alongside a wealth of information around parenting, childcare use and parental characteristics. In particular, as mothers participated in the home visits as the child’s primary caregiver, we collected rich data on maternal skills, including mothers’ years of education, verbal ability, IQ, depressive symptoms and knowledge of child development.⁷

Importantly, the household surveys also contain 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 home and about the types and frequencies of play activities the child engages in with an adult. Specifically, to measure play materials, we use questions about the number of different types of toys (e.g. toys to learn shapes, toys that induce physical movement) that the child has played with in the past 30 days. To measure play activities, we use questions about the activities performed by the primary caregiver or any other adult older than 15 with the child in the last 3 days. Such activities include, for example, reading or looking at picture books together, telling stories, and labelling things around the house.

The measures of child development that we collected in the home setting via maternal report include: language development (that is, the number of words and complex sentences the

⁷Specifically, as part of the household questionnaire, we administered on the mothers the Raven’s progressive matrices to test for IQ, the CES-D 10-item scale to test for depressive symptoms and the Knowledge of Infant Development Inventory (KIDI) to test for knowledge of child development).

child can say) using the vocabulary checklists in the Spanish Short-Forms of the MacArthur-Bates Communicative Development Inventories I and II (MacArthur); child temperament using Bates' Infant Characteristics Questionnaire (ICQ); and the attention focusing and inhibitory control scales in the short versions of the Early Children's Behavior Questionnaire (ECBQ).⁸ All of these were measured using age-appropriate items pre- and post-intervention, with the exception of the ECBQ which was administered at follow-up only. In addition to these assessments via maternal reports, we also had trained psychologists administer the Bayley Scales of Infant and Toddler Development III (Bayley) in community centres.⁹ These direct assessments of the child took place over an average period of 1.5 hours and were aimed at measuring children's cognitive, language and motor development in depth.

The measures of child development, maternal skills and parental investments that we use in this paper are described in detail in Appendix B. Appendix Table ?? reports the baseline characteristics of children, their mothers and their households. At baseline, the children are on average aged 18 months. About 10% of them were born premature and 14% of them were stunted. On average, their mothers are 26 years old, have about 7.5 years of education and 67% of them are either married or cohabiting. There were no compromises to the randomization protocol and hence there is no reason to believe there is any bias. Most baseline characteristics are very well balanced including the baseline skills of the children. Although the mean of a few characteristics is significantly different between treated and controls when tested individually (specifically among CESD scale items), none of these differences are significant at all when we allow for multiple hypothesis testing using the Romano and Wolf (2005) procedure.

⁸See Jackson-Maldonado et al. (2012) for MacArthur-Bates scales, Bates et al. (1979) for the ICQ and Putnam et al. (2006) for ECBQ.

⁹See Bayley (2006).

3.2 Factor Models and the measurement system

Our main aim is to interpret the experimental results within the context of a model of parental investments and human capital production functions. To fix ideas, suppose we wish to estimate a production function for child skills:

$$\theta_{t+1} = f_{t+1}(\theta_t, I_{t+1}, P, X_t) \tag{1}$$

where θ_t and θ_{t+1} are vectors of child’s skills at t and $t + 1$ respectively, I_{t+1} are parental investments that occur between the realizations of θ_t and θ_{t+1} , P are maternal skills and X_t is a vector of household characteristics, such as household composition. The production function allows us to understand the pathways through which the experiment might affect outcomes: changes in parental investments and/or changes in the production function $f(\cdot)$, reflecting, for example, better use of parental inputs.¹⁰

As Cunha and Heckman (2008) explain, an important obstacle to estimating such a function is that the skills and investments are inherently unobservable. The various measures described in Section 3.1 can be viewed as error ridden indicators for these underlying latent factors. Using any one set of these measures in place of the latent factors would lead to severely biased results, whether the model is linear or not. We thus follow the approach of Cunha et al. (2010)¹¹ and develop a measurement system linking the observed measures to latent factors and estimate the distribution of such factors.

Suppose we have \mathcal{M}_{kt}^θ measures of child’s skill θ_t^k of type k (e.g. cognitive or socio-emotional skills) in period t . Moreover, we also have \mathcal{M}_k^P measures of maternal skills P^k

¹⁰We use maternal skills as measured at baseline. However, we find no evidence of a treatment impact on any measures of cognitive skills or socio-emotional skills of the mother (the main primary caregiver in most households in our sample). This is in line with psychological evidence indicating that cognitive (as measured by IQ) is rank stable by the age of 10 (Almlund et al., 2011). While it is more plausible that the intervention could have changed maternal socio-emotional skills, we find no such evidence. Had these maternal measures changed they could have been an additional channel of impact.

¹¹More broadly this approach relates to the identification and estimation of nonlinear models with classical measurement error (Schennach, 2007; Shennach, 2004).

of type k . Finally, we have $\mathcal{M}_{\tau t}^I$ measures of parental investments I_t^τ of type τ (e.g. time or material investments) made between t and $t + 1$. We denote m_{kjt}^θ the j -th measure of child's skill of type k at t , m_{kj}^P the j -th measure of maternal skill of type k , and $m_{\tau jt}^I$ the j -th measure of parental investment of type τ at t . As we estimate a different joint distribution of latent factors for the control and treated groups, in what follows we index the measures and latent factors by the treatment subscript d , where $d = 0$ refers to the control group (no home visits) and $d = 1$ refers to the treatment group (some home visits).

As is common in the psychometric literature, we assume a dedicated measurement system, that is one in which each measure only proxies one factor (??). Although it is not necessary for identification, we maintain this assumption because it makes the interpretation of the latent factors more transparent and we find clear support for such a system in the data (see Appendix C). Assuming each measure is additively separable in the (log) of the latent factor it proxies,¹² we write the following system of equations mapping the j -th measure observed at some date t to the k -th latent (unobserved) factor for that date:

$$m_{kjd}^\theta = \mu_{kjt}^\theta + \alpha_{kjt}^\theta \ln \theta_{dt}^k + \epsilon_{kjt}^\theta \quad (2)$$

$$m_{kj}^P = \mu_{kj}^P + \alpha_{kj}^P \ln P_d^k + \epsilon_{kj}^P \quad (3)$$

$$m_{\tau jt}^I = \mu_{\tau jt}^I + \alpha_{\tau jt}^I \ln I_{dt}^\tau + \epsilon_{\tau jt}^I \quad (4)$$

where the terms μ_{kjt}^θ , μ_{kj}^P and $\mu_{\tau jt}^I$ are intercepts, the terms α_{kjt}^θ , α_{kj}^P and $\alpha_{\tau jt}^I$ are factor loadings, and the terms ϵ_{kjt}^θ , ϵ_{kj}^P and $\epsilon_{\tau jt}^I$ are measurement error mean zero terms which are assumed independent of the latent factors and of each other.¹³

¹²We specify the measurement equation such that measures proxy the log of a latent factor so that the latent factors only take positive values.

¹³The assumption that the errors are independent of each other can be relaxed somewhat. Some of the child cognitive outcomes, for example, 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 a different interviewer. However, it is certainly possible that measurement errors are correlated, even in this case from say child behavior, the implications of which should be studied in future research.

An important assumption we have made in writing the system above is that the measurement system is invariant between treated and controls. That is, the measurement system intercepts and factor loadings and the distribution of measurement errors are the same between treated and control groups. Importantly, this implies that any differences in the distribution of observed 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. It is possible, however, that the treatment changed the salience of some measures for the respective latent factor, which would be captured by a change in the factor loading. We investigate the validity of this assumption by estimating these parameters separately for treatment and control groups and testing if the factor loadings are different in the two groups. As we elaborate in Appendix D, we do not reject the hypothesis of equality between treatments. Going forward therefore, we maintain the assumption of measurement invariance as it restricts the number of free parameters and lead to improvements in efficiency.

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 on one of the measures (say the first) of each latent factor to 1, that is: $\alpha_{k1t}^\theta = \alpha_{k1}^P = \alpha_{\tau 1t}^I = 1, \forall t, \tau = \{M, T\}$ and $k = \{C, S\}$. When it comes to the child’s skills, we normalize the factor loading on the same measures at baseline and follow-up.¹⁴ We set the location of all the factors by fixing the mean of the latent factors in logs to 0 for the control group; the difference of the treatment location from that of the control (that is set to zero) is taken to be the average effect of the treatment.

With the assumptions and normalizations already made and based on the Kotlarski theorem and further extensions, Cunha et al. (2010) show that both the distribution of measurement errors and the latent factor distribution are non-parametrically identified so

¹⁴For cognitive skills, we normalize the factor loadings on the Bayley cognitive item both at baseline and follow-up. For socio-emotional skills, we normalize the factor loadings on the item measuring how difficult the child is in the ICQ.

long as we have at least three measures with nonzero factor loadings corresponding to each latent factor.¹⁵

While these assumptions are sufficient for identification, some of them could be relaxed as shown in the identification proofs in Cunha et al. (2010).¹⁶ 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 is one measure whose error is independent from those of other measures of the same factor.

An issue of practical importance relates to the scale of the latent factors and what they actually mean for measures of interest such as earnings. This is the issue of anchoring discussed in Cunha et al. (2010) who provide a theoretical treatment.¹⁷ In our paper we scale all cognitive measures for the children on the Bayley cognitive scale. This has a cardinal interpretation (the number of tasks completed correctly) and the same test is applied across different ages (up until 42 months), allowing for comparability.¹⁸ The lack of long-term longitudinal data prevents us from converting these units to future earnings or other adult outcomes of interest. For socio-emotional skills we also normalize to the same ICQ item (whether the child is difficult) in both periods. The estimates of the elasticity of substitution between inputs and of other production function parameters will depend on how inputs are anchored.

¹⁵See also Shennach (2004), Schennach (2007), Hu and Schennach (2008), Carneiro, Hansen, and Heckman (2003), Heckman, Pinto, and Savelyev (2013) and Cunha and Heckman (2008).

¹⁶See also Carneiro et al. (2003) and Cunha and Heckman (2008)

¹⁷Cunha et al. (2010) provide a general theoretical treatment of anchoring. In their main empirical results they anchor the measure of skills measured at the oldest age to years of education. They then assume that the same anchoring scale applies to measures of cognition and socio-emotional skill measured at earlier ages. This identifying assumption is probably unavoidable because skill at an early age is always going to be measured on some numerical scale without an immediate correspondence to say monetary units. Nielsen (2015) discusses using ordinal tests scores to measure achievement gaps.

¹⁸See Agostinelli and Wiswall (2016) on the importance of not rescaling over time.

3.3 Estimation of the measurement system and the latent factor distribution

All the necessary information to directly evaluate the impact of the intervention on the latent factors and to estimate the model of human capital formation is embodied in the joint distribution of the latent factors and some key demographic characteristics (such as the number of children). The specific estimation approach we follow is described in Attanasio et al. (2015) and we give a brief overview here.

We estimate two joint distributions: one for the treatment group and one for the control group. In each case, we approximate the joint distribution of log of the latent factors and the demographics as a mixture of two joint log-normal distributions. Allowing departure from the normal is crucial, since normality implies a linear conditional mean, which in our context would *impose* a Cobb-Douglas production function (linear in logs). We have experimented with more than two mixtures with no change in the results. In addition, we assume that the measurement errors are distributed as a joint normal distribution with means 0 and diagonal variance-covariance Σ_e , in keeping with the assumption that they are independent of each other.

A direct implication of the additive separability of the measurement equations (2) - (4), together with the assumption of that the factors follow a mixture of log-normal distributions and the assumption of normality of the additive measurement error, is that the joint distribution of measurements is given by a mixture of normals. This allows us to estimate the latent factor model with a simple two-step procedure. First, we estimate the parameters of the joint distribution of measurements and all other exogenous variables used in the model (such as demographics) by maximum likelihood, using the EM algorithm.¹⁹ Second, we 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. As we explain in Section 5.3, it is from these estimated

¹⁹See Arcidiacono and Jones (2003).

joint distributions of latent factors that we will draw treatment and control samples in order to both evaluate the intervention impacts and estimate our model of human capital accumulation.

3.4 Specification of the measurement system

An important preliminary step in implementing the measurement system above is to allocate measures observed in the data to particular factors, which is shown in Table 1.²⁰ The factor loading on the first measure is the one that is normalized to one and thus the one that defines the scale of the latent factor.

As reflected in the table, we did not necessarily use the same set of measures of the child’s skill at baseline and at follow-up, the main reason being that we only included age-appropriate items that provide relevant information about the latent skill. For example, the MacArthur item measuring the number of complex phrases a child can say is too advanced for children at 1-2 years old and hence was only administered at follow-up when children were between 2.5 to 3.5 years old. Similarly, with respect to socio-emotional skills, the ECBQ is designed to measure temperament among children aged 3-7 and therefore was only administered at follow-up.²¹ However, under the assumptions made for the measurement system, this does not pose a problem for the approach we follow. Importantly, in both rounds, we use the same measure to normalize the child’s baseline cognitive and socio-emotional skills, so magnitudes are comparable over time.

In our model we use mother’s skills to control for parental background. During the data collection process, we had to focus only on the mother’s skills (who is almost always the principal caregiver and often a single mother) because of resource constraints and in order to keep interview times at a reasonable level. In so doing, it is possible that we miss the

²⁰We perform an exploratory factor analysis reported in Appendix C to identify in a preliminary step the relevant measures and their allocation to factors.

²¹The ICQ is in principle designed for children up to 2 years old. We administered the same questions of the ICQ at baseline and follow-up after consultation with the developer of the test.

Table 1: Measurement system

Latent factor	Measurement	Survey	% Signal	
			Controls	Treated
Child’s cognitive skills at FU (θ_{t+1}^C)	Bayley: cognitive	FU	0.78	0.79
	Bayley: receptive language	FU	0.75	0.75
	Bayley: expressive language	FU	0.78	0.78
	Bayley: fine motor	FU	0.60	0.60
	MacArthur: words the child can say	FU	0.63	0.63
	MacArthur: complex phrases the child can say	FU	0.51	0.51
Child’s cognitive skills at BA (θ_t^C)	Bayley: cognitive	BA	0.73	0.69
	Bayley: receptive language	BA	0.75	0.71
	Bayley: expressive language	BA	0.77	0.73
	Bayley: fine motor	BA	0.63	0.58
	MacArthur: words the child can say	BA	0.49	0.44
Child’s socio-emotional skills at FU (θ_{t+1}^S)	ICQ: difficult (-)	FU	0.74	0.70
	ICQ: unsociable (-)	FU	0.35	0.30
	ICQ: unstoppable (-)	FU	0.60	0.54
	ECBQ: inhibitory control	FU	0.72	0.68
	ECBQ: attention focusing	FU	0.29	0.25
Child’s socio-emotional skills at BA (θ_t^S)	ICQ: difficult (-)	BA	0.61	0.67
	ICQ: unsociable (-)	BA	0.27	0.33
	ICQ: unadaptable (-)	BA	0.31	0.37
	ICQ: unstoppable (-)	BA	0.18	0.22
Material investment at FU (I_t^M)	FCI: Number of different types of play materials	FU	0.95	0.97
	FCI: Number of books to paint and draw	FU	0.14	0.23
	FCI: Number of toys to learn movement	FU	0.62	0.75
	FCI: Number of toys to learn shapes	FU	0.69	0.80
	FCI: Number of toys bought	FU	0.64	0.77
Time investment at FU (I_t^T)	FCI: Number of different types of play activities	FU	0.88	0.93
	FCI: Times told a story to child in last 3 days	FU	0.67	0.80
	FCI: Times read to child in last 3 days	FU	0.75	0.86
	FCI: Times played with toys and the child in last 3 days	FU	0.59	0.74
	FCI: Times labelled things to child in last 3 days	FU	0.60	0.75
Mother’s cognitive skills at BA (P^C)	Mothers’ years of education	BA	0.56	0.53
	Mother’s vocabulary	FU	0.64	0.62
	Mother’s Raven’s score (IQ)	FU2	0.51	0.48
	FCI: No. of books for adults at home	BA	0.37	0.35
	FCI: No. magazines and newspapers at home	BA	0.19	0.18
Mother’s socio-emotional skills at BA (P^S)	CESD: Did you feel depressed? (-)	BA	0.55	0.68
	CESD: Are you bothered by what usually don’t? (-)	BA	0.23	0.33
	CESD: Did you have trouble keeping mind on doing? (-)	BA	0.34	0.47
	CESD: Did you feel everything you did was an effort? (-)	BA	0.34	0.47
	CESD: Did you feel fearful? (-)	BA	0.41	0.54
	CESD: Was your sleep restless? (-)	BA	0.25	0.35
	CESD: Did you feel happy?	BA	0.29	0.41
	CESD: How often did you feel lonely last week? (-)	BA	0.36	0.49
	CESD: Did you feel you couldn’t get going? (-)	BA	0.39	0.52

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. “BA” refers to Baseline, “FU” refers to the first-follow-up survey and “FU2” refers to the second follow-up survey collected 2 years after the intervention ended. The symbol (-) indicates that the scoring on these measures was reversed so that a higher score on the corresponding latent factor means a higher level of skill.

genetic or other influence of the father; however, we expect to be capturing at least some of that by conditioning on the baseline skills of the child. We use baseline measures to extract two factors measuring the mother’s cognitive and socio-emotional skills, with the exception of the vocabulary test, which was administered at follow up and the Raven’s score which was administered at a later round of data collection (2 years after the end of the intervention). In both cases we checked and the intervention had not impact on the scores.

The parameters of the measurement system for treatment and control are estimated together with the latent factor distributions as described above. We report estimates of the factor loadings and of distribution of measurement errors in Appendix C. To assess the extent of information relative to measurement error contained in each of the measures, we compute the signal-to-noise ratio measuring the fraction of the variance of each measure driven by signal. For example, for the j -th measure of child’s skills of type k , this ratio is defined as:

$$s_j^{\ln \theta^k} = \frac{(\alpha_j^k)^2 \text{Var}(\ln \theta^k)}{(\alpha_j^k)^2 \text{Var}(\ln \theta^k) + \text{Var}(\epsilon_j^k)}$$

where we have assumed that the j -th measure of latent factor θ^k can be written, simplifying notation, as:

$$m_j^\theta = \mu_j^k + \alpha_j^k \ln \theta^k + \epsilon_j^k$$

The last two columns of Table 1 report the signal-to-noise ratio for each of the measures used in the analysis for the control and treated groups separately. These numbers can be different because the joint distribution of latent factors is allowed to be different between the two groups. Clearly, there is much variation in the amount of information contained in each measure of the same factor. For example, 78% of the variance in the *Bayley: Cognitive* item is due to signal, whereas only 51% of the variance in the *Mac Arthur: Complex Phrases* item is due to signal. Overall, most measures are far from having 100% of their variance accounted

for by signal, which emphasizes the importance of accounting for measurement error through the latent factor model. As we discuss in Section 4, the presence of pervasive measurement error implies that some measures may register a significant impact of the program and others may not. Moreover, it also illustrates the usefulness of the latent factor approach in modeling human capital accumulation and parental investments: without such an approach, one would risk to obtain severely attenuated coefficients, masking the importance of investments and background variables on child development.

4 Short-term impacts of the intervention

In this section, we document the impacts of the intervention on child’s cognitive and socio-emotional development as well as parental investments, observed at first follow-up, just after the 18 month-long intervention ended. In each panel of Table 2, we report the estimated impacts of receiving the home visits on one of four sets of outcomes: (i) cognitive development; (ii) socio-emotional development; (iii) parental investment in play materials; (iv) parental investment in play activities. In addition to the impact on each measure, we also report the impact on the mean of the corresponding log latent factor. As mentioned earlier, we focus on the impact of the psychosocial stimulation component of our intervention because there were no significant impact of micro-nutrient supplementation on any child developmental outcomes (Attanasio et al., 2014). If we explicitly control for the fact that half the stimulation group also received micronutrient supplementation, the impact on cognition and receptive language remains virtually the same, with a very small increase in the point estimates we report below (see Appendix Table ??).

4.1 Impacts on child development

The top panel of Table 2 summarizes the short-term impact of the intervention. These imply an increase of 0.25 of a standard deviation (SD) in cognitive development and an

Table 2: Treatment impacts on raw measures and factors

	Treatment effect	
	Point estimate	Stand error
<i>A - Child's cognitive skills at follow-up</i>		
Bayley: cognitive	0.250	(0.063)
Bayley: receptive language	0.175	(0.063)
Bayley: expressive language	0.032	(0.062)
Bayley: fine motor	0.072	(0.060)
MacArthur: words the child can say	0.092	(0.064)
MacArthur: complex phrases the child can say	0.058	(0.054)
Cognitive factor	0.115	(0.054)
<i>B - Child's socio-emotional skills at follow-up</i>		
ICQ: difficult	-0.074	(0.045)
ICQ: unsociable	-0.041	(0.054)
ICQ: unstoppable	-0.032	(0.054)
ECBQ inhibitory control	-0.003	(0.058)
ECBQ: attention focusing	0.067	(0.049)
Socio-emotional factor	0.086	(0.043)
<i>C - Material investment at follow-up</i>		
FCI: Number of different types of play materials	0.215	(0.064)
FCI: Number of books to paint and draw	-0.133	(0.056)
FCI: Number of toys to learn movement	-0.048	(0.065)
FCI: Number of toys to learn shapes	0.416	(0.088)
FCI: Number of toys bought	0.024	(0.061)
Material investment factor	0.225	(0.071)
<i>D - Time investment at follow-up</i>		
FCI: Number of different types of play activities	0.277	(0.050)
FCI: Times told a story to child in last 3 days	0.138	(0.060)
FCI: Times read to child in last 3 days	0.362	(0.062)
FCI: Times played with toys and the child in last 3 days	0.175	(0.059)
FCI: Times labelled things to child in last 3 days	0.137	(0.048)
Time investment factor	0.298	(0.072)

Note: All scores have been internally standardized non-parametrically for age and are expressed in standard deviation units (see Appendix B for details about the measures and the standardization procedure). The effects relating to the *latent factors* are in log points. Coefficients and standard errors clustered at the municipality level (in parentheses) from a regression of the dependent variable measured at follow-up on an indicator for whether the child received any psychosocial stimulation and controlling for child's sex; tester effects and baseline level of the outcome. Sample size: 96 communities and individual observations range from 1,262 to 1,326 depending on the outcome.

increase of 0.175 SD in receptive language, assessed using the Bayley.²² The cognitive factor summarizing all these effects shows a substantial and significant increase of 11% (0.11 log points) amongst the treated group relative to the control group.

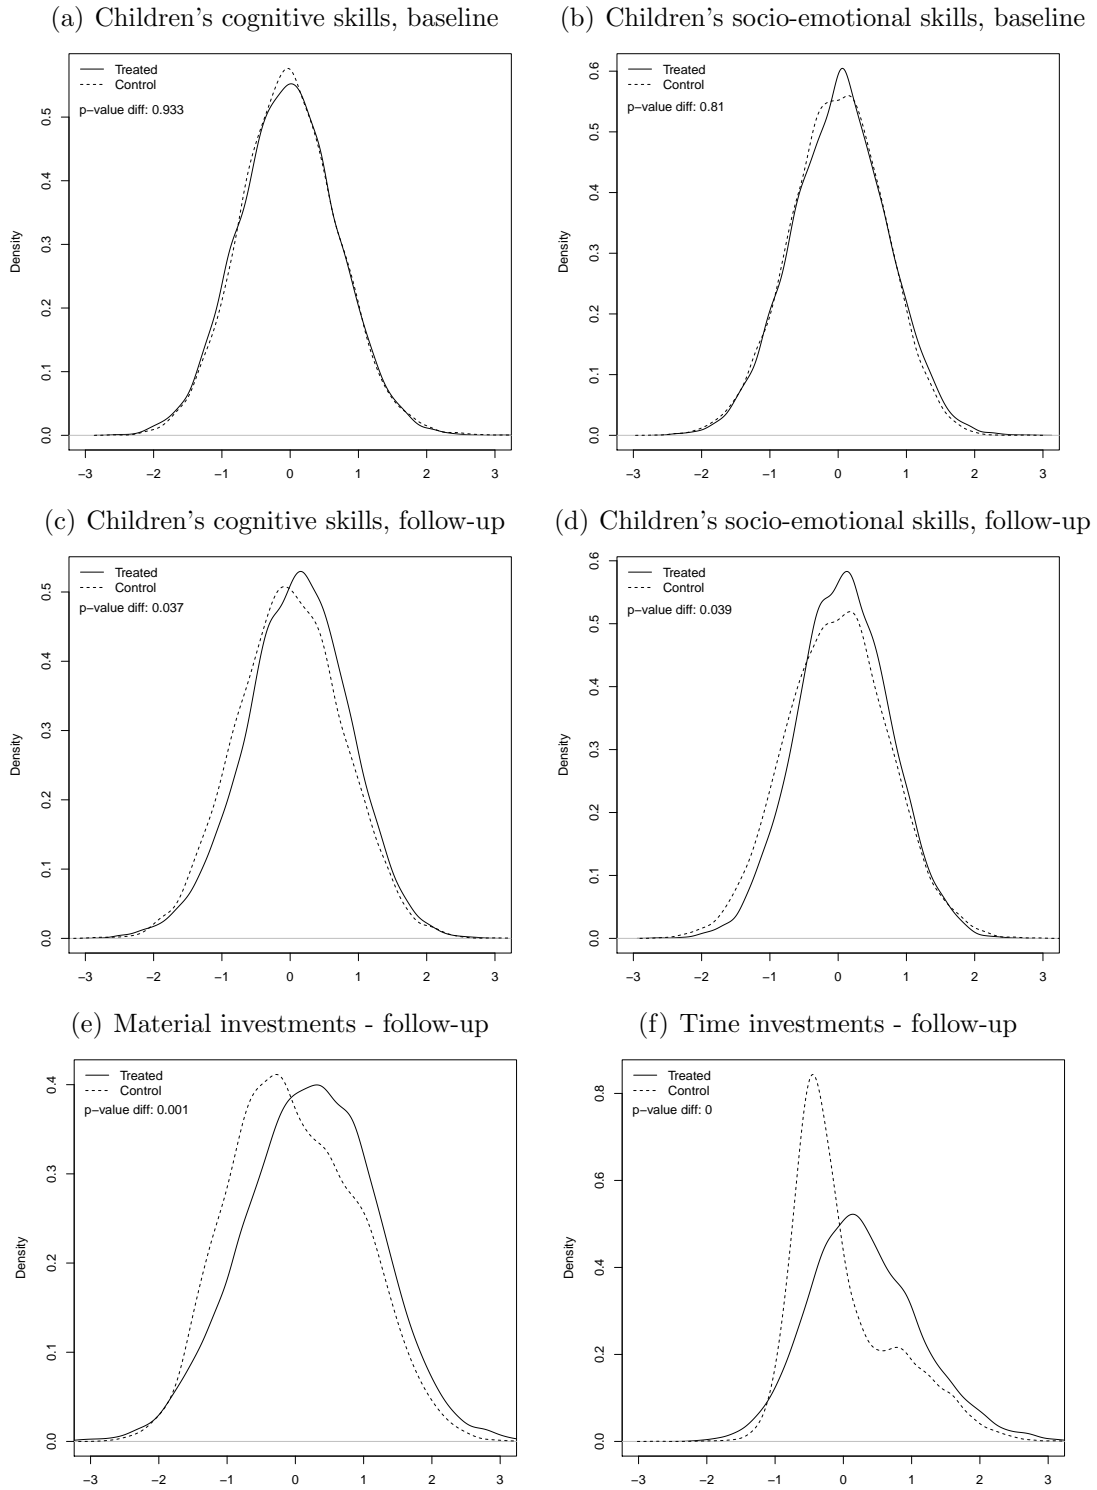
The second panel of the table also shows that the intervention led to an improvement in some dimensions of socio-emotional development. In particular, it resulted in a 0.07 SD decrease in the dimension of the ICQ scale measuring difficult behavior; none of the other three components of the ICQ scale were individually significant. However, there is a significant improvement (p-value<0.05) in the socio-emotional skills factor, which increased by 8.6%. These results illustrate the power of the latent factor approach in extracting information from the various measures.

Having identified and estimated the entire distribution of factors for each treatment and control group, we can study whether the intervention 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 for the control and treated groups and perform a Kolmogorov-Smirnov (K-S) test of the hypothesis that the corresponding CDFs are equal to each other (the p-values of the test are reported in the figure and have been derived using the bootstrap). The first two panels show the distribution, in treatment and control villages, of cognitive and socio-emotional skills at baseline. The two densities overlap each other and the K-S test cannot reject that they are equal to each other, thus confirming that our sample is balanced.

The following two panels depict the distribution of cognitive and socio-emotional factors at follow-up. In the case of the cognitive factor, we see that the shift in the mean reported in Table 2 reflects a shift in the entire distribution. For the socio-emotional factor, however, the shift occurs mainly for children below the median. For both types of skills, the K-S test rejects the equality of the two distributions at the 5% significance level (p-values are 0.037

²²These 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.

Figure 1: Kernel densities of latent factors



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. For each factor, we perform a Kolmogorov-Smirnov test using the bootstrap and accounting for the entire estimation procedure. p-values reported in each panel.

for cognitive skills and 0.039 for socio-emotional skills).

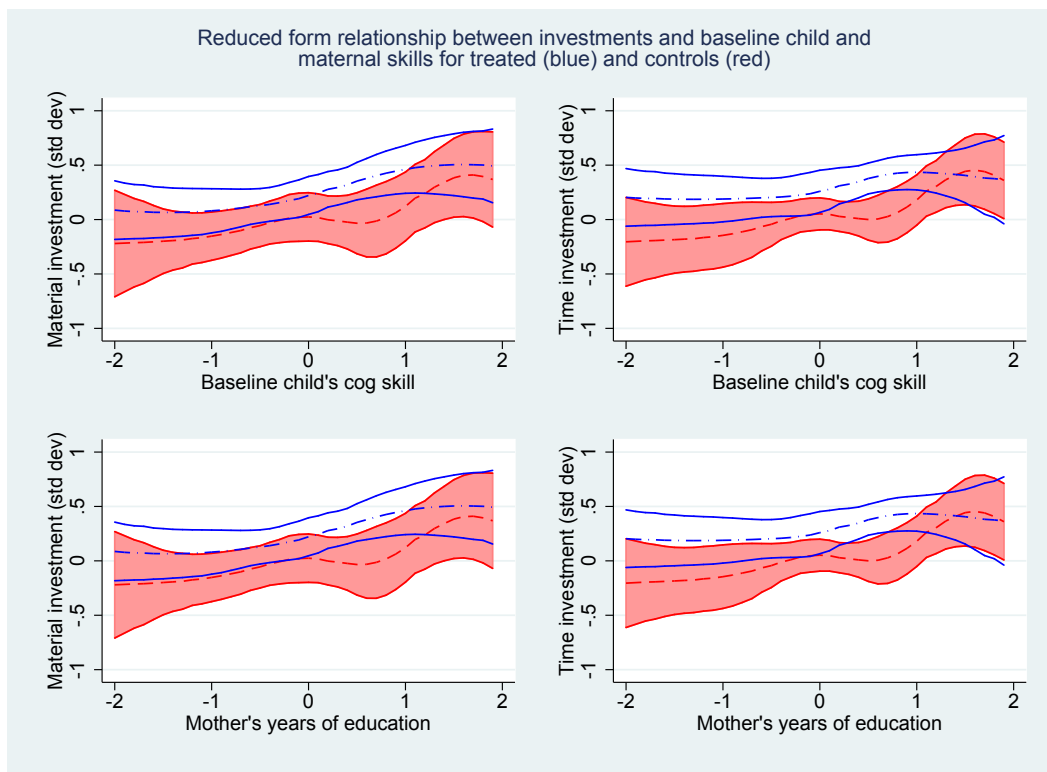
4.2 Suggestive evidence of mechanisms

The bottom two panels of Table 2 report the impacts of the stimulation intervention on parental investment. We observe substantial impacts on several individual items, as measured by the FCI, as well as the two investment latent factors. Among materials the increase is not uniform. Specifically, there is an increase for some toys but a reduction in coloring books; the overall factor registers an increase, which is highly significant. The lower panel shows that all types of time activities increase. It is quite clear from these results that parents are doing much more as a result of the intervention. This is particularly important, given that the measured activities are separate from the intervention itself. In the last two panels of Figure 1, we notice a strong shift to the right of the distributions of both the material and time investment factors. For either type of investments, the K-S strongly rejects the equality of the corresponding CDFs between control and treated groups.

In Figure 2 we take a specific measure of material and time investment and show how these vary with treatment across the distribution of child’s baseline Bayley cognitive score and mother’s year of education (both in standard deviation units). These graphs show first that the intervention increased investments more or less by the same amount at all levels of both mother and child’s baseline skills. Moreover, investments increase with both of these baseline skills. This suggests that investments vary with parental background and that parents provide greater attention to higher ability children because they may perceive a complementarity between investments and initial skills. These are themes we will return to within our model.

The impacts we observe on parental investments are suggestive that one mechanism through which the intervention might have improved child development was by promoting parental investments in children. In the general production function set out in equation 1 above, this would correspond to a shift in the distribution of I_t . There are however other

Figure 2: Investments by treatment and baseline skills of the mother and the child



Note: Predicted values and 95% confidence intervals produced using cubic splines.

mechanisms through which the program could have been effective. For example, the weekly home visits could have had a direct impact on skills and/or could have improved the quality or effectiveness of parental investments, thus changing the structure of the production function in addition to changing the distribution of its inputs. To investigate these hypotheses further, we specify a model of parental investments and human capital production functions that define the relationship between child skills at follow-up and parental investments (amongst other inputs). Within that framework, we then test whether the intervention changed the production function parameters.

5 The accumulation of human capital in the early years and the role of the intervention

We specify a model of parental investments and child skill formation. Parents can choose to invest in materials and time, and skills take two dimensions, namely cognitive and socio-emotional skills. We refer to the baseline period as t , when children were between 12 to 24 months old, and to the post-intervention period as $t + 1$, when children were between 30 to 42 months old.

Child skills at $t + 1$ are assumed to be a function of the vector of child skills at t , maternal skills at t , parental investments in the intervening period and random shocks. We think of investments as parental choices. However, rather than modeling investment choices resulting from the dynamic optimization of a household problem as in DelBoca, Flinn, and Wiswall (2014), we estimate a pair of reduced form investment equations, which could be interpreted as an approximation to those derived (numerically) in a full structural model. By not imposing any restrictions from a structural model we do not have to take a stance on whether parents now the process of child development reflected in the structure of the production function.

Beyond the information on the process of child development that the model yields, it offers a framework to understand the mechanisms underlying the intervention. One mechanism could be an increase in parental investments. Indeed, the intervention aims to strengthen child-mother interactions and encourage mothers to engage more with the child by incorporating age-appropriate play activities in the daily routine, introducing new toys constructed with home-made materials and spending time reading, telling stories or singing. However, it is also possible that investments could decline as parents shift their attention and resources elsewhere (for example, to other children) because they perceive the intervention itself as an in-kind transfer and thus re-optimize the allocation of resources within the household. Such

crowding-out of private resources is a standard concern in programs that target children, but crowding-in is also a plausible possibility (Gelber and Isen, 2010; Jacoby, 2002), particularly in the presence of liquidity constraints and/or if lumpy investments are required to obtain better life outcomes for their children. Indeed, if the in-kind transfer partially bridges the gap needed for such better outcomes and if parents now perceive the previously unattainable goal of a better outcome for their child as feasible, they are likely to increase their investment. Effectively, a small addition may have a large return relative to the case where no intervention was taking place.

A second mechanism through which the intervention may have operated is by changing the production function itself. On the one hand, the stimulation provided during the home visits may be a new input in the development of the child, and this would be captured by a shift in the TFP or other parameters of the production function. On the other hand, parents, now guided by the intervention, may use time and resources in a more effective way. This interpretation implies that, despite the richness of our data, some aspects of investment quality are not captured by our measures and thus get embodied in the estimates of the production function.²³

Finally, the intervention could also have affected maternal cognitive or, more plausibly, socio-emotional skills. Many of the mothers (37%) were depressed at baseline according to the CESD scale and it is plausible that the treatment mitigated this. Although we checked for such impacts, we did not detect any differences in our measures of maternal skills (either cognitive or socio-emotional skills) between control and treated after the intervention; thus this potential change is not a mechanism that contributed to the outcome. In our estimated model we only include baseline maternal skills.²⁴

²³We made every effort to collect both time and resource use carefully targeted to the child with an emphasis on items that can drive development. For example, one of our measures is the number of times spent reading with the child in the last 3 days. Yet, it is still a possibility that as a result of the intervention, parents may be more able to select age-appropriate or stimulating stories to read with their child. Our measure of the frequency with which parents read with their child would not pick up this change in the quality of interaction, which would instead be picked up by a shift in the productivity of time investments.

²⁴The effect of the intervention on the principal component factor of the CES-D scale items at follow-

5.1 The production function for human capital

In what follows we introduce the subscript i for individual. The subscript t indicates baseline period and $t + 1$ follow up. Investments, which carry a $t + 1$ subscript, are measured at follow up and refer to the period between the two surveys. Finally, any element of the model below (parameters or variables) that could have been affected by the intervention includes a subscript d , where $d = 1$ denotes intervention and $d = 0$ control.

We assume that the stock of skills in period $t + 1$ is determined by the vector of child's baseline cognitive and socio-emotional skills θ_{it} embodying the initial conditions at the time of observation (possibly including any paternal influence), the mother's cognitive and socio-emotional skills denoted by P_{it}^C and P_{it}^S respectively, and the investments I_{it+1} made by the parents between t and $t + 1$. We also allow for the effect of a variable η_{it+1}^k that reflects unobserved shocks or omitted inputs. As with skills, parental investments I_{it+1} can be a multi-dimensional vector. Here, we distinguish between material and time investments, which we denote as I_{it+1}^M and I_{it+1}^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:²⁵

$$\begin{aligned} \theta_{idt+1}^k = & A_d^k [\gamma_{1d}^k (\theta_{it}^C)^{\rho_k} + \gamma_{2d}^k (\theta_{it}^S)^{\rho_k} + \gamma_{3d}^k (P_{it}^C)^{\rho_k} + \gamma_{4d}^k (P_{it}^S)^{\rho_k} \\ & + \gamma_{5d}^k (I_{idt+1}^M)^{\rho_k} + \gamma_{6d}^k (I_{idt+1}^T)^{\rho_k} + \gamma_{7d}^k n_{it}^{\rho_k}]^{\frac{1}{\rho_k}} e^{\eta_{it+1}^k} \quad k \in \{C, S\} \end{aligned} \quad (5)$$

where n_{it} is the number of children in the household. This is to allow for the possibility that the presence of siblings affects child development because of spillover effects and more broadly because of the learning and socialization that can be achieved by interacting with other older children.²⁶ It is possible, on the other hand, that the presence of siblings dilute

up is 0.13 of a standard deviation (with a p-value of .12), which is indicative of an improvement but too insignificant to rely upon.

²⁵Cunha et al. (2010) also use a CES, while Cunha and Heckman (2008) use a log linear production function.

²⁶Since our subject children are 12-24 months old at baseline these are almost always older children.

attention and resources, but this will be captured by the investment functions as we explain below.

A_d^k is a factor-neutral productivity parameter and depends on the treatment status of the child (d) to capture the potential direct effect of the home-visitor stimulating the child during her weekly visit. $\rho_k \in (-\infty, 1]$ determines the elasticity of substitution, given by $1/(1 - \rho_k)$, between the inputs affecting the accumulation of skill k . Under such parameterization, as $\rho_k \rightarrow -\infty$, the inputs become perfect complements. As $\rho_k \rightarrow 1$, the inputs become perfect substitutes. The intervention could, in principle, affect any of the parameters of the production function and, as we discuss later, we test empirically whether this is the case.

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 cognitive skills and that of 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 and translog production functions. We could not reject the CES functional form however and so we maintain this functional form assumption throughout the application.

5.2 Parental investments

We model investments as a function of the child and the mother's baseline skills and the number of children in the household.²⁷ The number of children in the household may dilute both the resources and the time devoted to our subject child. We also include a vector of

²⁷We use the total number of children so the minimum is 1.

variables Z_{it} discussed below. The estimating investment equations are

$$\begin{aligned} \ln(I_{idt+1}^\tau) = & \lambda_{0d}^\tau + \lambda_{1d}^\tau \ln(\theta_{it}^C) + \lambda_{2d}^\tau \ln(\theta_{it}^N) + \lambda_{3d}^\tau \ln(P_i^C) + \lambda_{4d}^\tau \ln(P_i^S) \\ & + \lambda_{5d}^\tau \ln(n_{it}) + \lambda_{6d}^\tau \ln(Z_{it}) + u_{it+1}^\tau, \end{aligned} \quad \tau = \{M, T\} \quad (6)$$

As implied by the subscript d all coefficients could change with the treatment. This is a hypothesis we test. The effect of background variables on parental investment, given child initial conditions, is an important potential source of socio-economic gradients in child development. Moreover, the extent to which investments increase with child initial abilities is a reflection of parental beliefs about the heterogeneity of returns to such investments as well as parental taste for redistribution among children.

Parental investments are an input in the production function. However, they may be endogenous, i.e. it may be that $E(\eta_{it+1}^k | I_{it+1}) \neq 0$. In particular 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 η_{it+1}^k biasing downwards the impact of investments. 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.

To explain how we deal with investment endogeneity in our context, we rewrite the

production function in logs :

$$\begin{aligned} \ln(\theta_{idt+1}^k) = & \frac{1}{\rho_k} \ln[\gamma_{1d}^k(\theta_{it}^C)^{\rho_k} + \gamma_{2d}^k(\theta_{it}^S)^{\rho_k} + \gamma_{3d}^k(P_{it}^C)^{\rho_k} + \gamma_{4d}^k(P_{it}^S)^{\rho_k} + \gamma_{5d}^k(I_{idt+1}^M)^{\rho_k} \\ & + \gamma_{6d}^k(I_{idt+1}^T)^{\rho_k} + \gamma_{7d}^k n_{it}^{\rho_k}] + \ln(A_d^k) + \eta_{it+1}^k, \quad k = \{C, S\} \end{aligned} \quad (7)$$

Endogeneity of investment implies correlation of the error components u_{it+1}^T with each of the η_{it+1}^k . Using a control function approach, we assume that:

$$E(\eta_{it+1}^k | I_{idt+1}^M, I_{idt+1}^T, Q_{idt}) = \delta_{1d}^k u_{idt+1}^M + \delta_{2d}^k u_{idt+1}^T \quad (8)$$

where Q_{idt} represents all other inputs into the production function and the two investment equations, including the variables Z_{it} and treatment.

Once the parameters of the investment functions are estimated, we recover \hat{u}_{it+1}^T and \hat{u}_{it+1}^M , the estimated residuals from the time and material investment equations respectively (6), which we include as regressors when estimating the production functions:

$$\begin{aligned} \ln(\theta_{idt+1}^k) = & \frac{1}{\rho_k} \ln[\gamma_{1d}^k(\theta_{it}^C)^{\rho_k} + \gamma_{2d}^k(\theta_{it}^S)^{\rho_k} + \gamma_{3d}^k(P_{it}^C)^{\rho_k} + \gamma_{4d}^k(P_{it}^S)^{\rho_k} + \gamma_{5d}^k(I_{idt+1}^M)^{\rho_k} \\ & + \gamma_{6d}^k(I_{idt+1}^T)^{\rho_k} + \gamma_{7d}^k n_{it}^{\rho_k}] + \ln(A_d^k) + \phi_M^k \hat{u}_{it+1}^M + \phi_T^k \hat{u}_{it+1}^T + v_{it+1}^k, \quad k = C, N \end{aligned} \quad (9)$$

For identification, our approach requires that we have at least as many exclusion restrictions Z as endogenous variables. That is, it requires that we have at least 2 variables that determine investment choices but do not enter the production function directly and that are uncorrelated with the production function shocks, as is implicit in the control function assumption of equation (8). A natural candidate for Z would be the intervention we described above, as it was allocated randomly across villages. However, we have already argued that the intervention can change the production function, a hypothesis which we test; so initially we do not use the intervention as an excluded instrument, although we consider this later in the paper. In any case there are two investments and only one intervention so more

instruments are needed.

Economic theory suggests that variables that exogenously shift the household’s resources might be valid instruments, since they impact parental investment decisions through the budget constraint without entering directly the production function. In this spirit, we use average male wages in the child’s town/village, the distance of the community from the regional capital and an indicator for whether the mother is married (which includes cohabiting) as variables that determine resources but do not enter the production function directly.

The validity of aggregate community level wages rests on the assumption that wage differences across communities only stem from differences in demand for labor. They could also be attributable to differences in individual preferences as long as these are not related in any way to the unobservables in the production function. Similarly, the distance variable relies on the assumption that proximity to the regional capital makes investments cheaper but does not affect child development *per se* nor correlates with omitted inputs (given those already included).

Perhaps the most controversial instrument is marital status (measured at baseline), which we exclude from the production function. The caveat is that the presence of a father figure (and his characteristics) may have a direct impact on child development. Hence the assumption that this variable is excludable from the production function implies that the role of a father figure, if any, works exclusively through the initial conditions of the child, the mother’s ability (through sorting) and the budget constraint. In what follows, we investigate the validity of these exclusion restrictions in various robustness exercises.

5.3 Estimation of the investment and production functions

In section 3.3, we discussed the estimation of the joint distributions of all latent factors and of the demographic variables Z , which are used as instruments or conditioning variables. We estimate one such distribution for the treatments and one for the controls. These distributions contains all the relevant information in the data for estimating the investment and

production functions and have been approximated each by a mixture of two normals.

To estimate the investment and production functions we draw a synthetic dataset of 10,000 observations from the treatment distribution and 10,000 from the control distribution. Using these simulated datasets, we first estimate the log-linear investment functions by ordinary least squares and construct the residuals \hat{u}_{idt+1}^τ ($\tau \in \{M, T\}$) that serve as control functions. 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. We compute standard errors and confidence intervals using the bootstrap.²⁸

6 Results

6.1 Estimates of the investment functions

We present estimates of the investment equations in Table 3. The first column presents the equation for material investments and the second column for time investments.

The first striking result is the impact of treatment on investments: it increases materials by 32% and time by 37%, and both effects are highly significant. 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 in log-investments seems to have been uniform across groups with differing backgrounds, implying an equal proportionate increase in investments.²⁹

Thus, the intervention increased the time and the resources that parents provide to children. Referring back to the measurement system (Table 2), it is worth noting that the time

²⁸We draw 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 starting with the joint distributions described above, for each one of the pseudo-sample. For each of the parameters, we then compute the standard deviation of its distribution based, along with various percentiles to compute the corresponding confidence intervals.

²⁹The estimates where all parameters of the investment functions are allowed to vary with treatment are shown in Web Appendix Table ???. 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: the p -value for the material investment equation is 0.577 and the p -value for the time investment equation is 0.667.

Table 3: Material and time investment equations

Dependent variable:	Log Material investment	Log Time Investment
Constant	-0.001 [-0.053, 0.046]	0.007 [-0.043, 0.042]
Treat	0.315 [0.158, 0.464]	0.385 [0.214, 0.542]
Log child's cognitive skill (t)	0.120 [-0.004, 0.21]	0.071 [-0.029, 0.164]
Log child's socio-emotional skill (t)	-0.020 [-0.126, 0.105]	0.000 [-0.103, 0.144]
Log mother's cognitive skill	0.787 [0.591, 0.988]	0.357 [0.089, 0.497]
Log mother's socio-emotional skill	0.082 [-0.023, 0.158]	0.100 [0.004, 0.175]
Log number of children (t)	-0.031 [-0.1, 0.044]	-0.121 [-0.194, -0.059]
Married (t)	0.122 [0.065, 0.173]	0.107 [0.047, 0.169]
Log avg males wages (t+1)	0.090 [-0.012, 0.155]	-0.019 [-0.113, 0.06]
Distance to region's capital	-0.099 [-0.204, -0.038]	-0.112 [-0.206, -0.038]
F-test (<i>p-value</i>)		
Male wage, distance, married	0.019	0.026
Male wage, distance	0.030	0.049
Cragg-Donald rank test (<i>p-value</i>)		
Male wage, distance, married		0.078
Male wage, distance		0.054

Note: Estimates based on 10,000 draws from estimated joint distribution of factors and instruments. Bootstrap 95% confidence intervals in square brackets based on 1,000 replications.

inputs are measured in a way that target 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 the child's cognitive skills affect material investment, but her socio-emotional skills have no impact on either type of investment. Child's cognitive skills have a small impact on time investments, but it is only significant at the 10% level. 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 only affect time investments significantly.

The implication of these results is that, even at this early stage in life and within this relatively deprived population, we can trace some of the origins of future inequality: higher human capital families invest more in their children, even *conditional* on child initial cognition and, moreover, children with higher early cognitive development obtain themselves more resources from their parents. Everything, in other words, pushes towards perpetuating and reinforcing initial inequalities.

Married mothers invest more time and more materials, but the overall number of children at baseline reduces time investments. We also 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. Wages have no significant effect on time investments. We also find that investments increase when the village is closer to the regional capital, reflecting the better access to services and overall lower costs: a 10% increase in distance reduces material and time investments by about 1%.

Of the above variables, marital status and average male wages in the village are excluded from the production function and serve as instruments. These together with the distance from the regional capital are jointly significant in each of the investment equations with p-values of 1.9% for material and 2.6% for time. However, there may be some cause for concern about the strength of the instruments in the two equation system since the p-value of the Cragg-Donald rank test is 7.8% if we consider all three instruments and 5.4% if we consider just male wages and distance to regional capital. However, the joint significance of the two

control functions should still provide a good indication of whether we can take investments as exogenous, given that the instruments are significant in the two investment equations. We also experiment with more parsimonious specifications of the production functions as we describe below.

6.2 Estimates of the production functions

For each of the two skills, we report estimates of the parameters of the production function in Table 4. The first two columns exclude the control functions (treating investments as conditionally exogenous) and the second includes them.

We start by considering whether the production function changes as a result of the intervention. This could happen for a number of reasons. First, the weekly session of the home visitor with the child can be thought of as a new input; second the intervention could lead to a better use of measured inputs by parents or equivalently an improvement in the unmeasured quality of these inputs. These are possible channels through which the intervention could affect outcomes over and above inducing more investments through its emphasis on parenting and the direct involvement of the mother in the home visit.

While TFP increased for the production function of cognitive skills when investments are treated as exogenous, in all other cases (including socio-emotional skills with exogenous investments) TFP does not change significantly. Moreover, there is no evidence that the remaining parameters changed.³⁰

Turning now to whether investments should be taken as endogenous, the coefficients on the residuals for material and time investment (the control functions) reported in the first two columns of Table 4 are not significant except in the case of material investments for socio-emotional skills (marginally). The joint test that the control functions are insignificant

³⁰The tests that the coefficients are the same across intervention and control groups, other than TFP A , have p -values of 0.25 for cognitive skills and 0.061 for socio-emotional skills when assuming investments are exogenous and 0.879 and 0.889 respectively when allowing for endogeneity. The unrestricted production functions are shown in the Appendix Table ??

Table 4: Production functions for cognitive and socio-emotional skills

	<i>Without control function</i>		<i>With control function</i>	
	<i>Cognitive skill</i>	<i>Socio-emotional skill</i>	<i>Cognitive skill</i>	<i>Socio-emotional skill</i>
TFP	0.985	1.004	0.998	1.006
	[0.969, 1.011]	[0.976, 1.044]	[0.971, 1.013]	[0.964, 1.039]
TFP \times Treatment	0.097	-0.02	0.102	0.061
	[0.016, 0.216]	[-0.091, 0.071]	[-0.006, 0.289]	[-0.055, 0.192]
Baseline cognitive skill	0.629	0.095	0.629	0.145
	[0.562, 0.744]	[0.013, 0.178]	[0.539, 0.777]	[0.033, 0.279]
Baseline socio-emotional skill	0.03	0.467	0.046	0.452
	[-0.066, 0.117]	[0.373, 0.615]	[-0.056, 0.142]	[0.357, 0.611]
Maternal cognitive skill	0.197	0.01	0.105	0.278
	[0.078, 0.294]	[-0.133, 0.12]	[-0.135, 0.406]	[-0.014, 0.54]
Maternal socio-emotional skill	0.032	0.016	0.057	0.041
	[-0.024, 0.104]	[-0.059, 0.094]	[-0.013, 0.161]	[-0.042, 0.131]
Material investment	0.087	0.14	0.339	-0.224
	[0.023, 0.159]	[0.059, 0.234]	[-0.108, 0.669]	[-0.657, 0.177]
Time investment	0.000	0.14	-0.206	0.179
	[-0.093, 0.078]	[0.042, 0.229]	[-0.523, 0.07]	[-0.218, 0.479]
Number of children	0.025	0.132	0.030	0.128
	[-0.027, 0.063]	[0.077, 0.17]	[-0.029, 0.072]	[0.072, 0.174]
Material investment residual			-0.268	0.381
			[-0.616, 0.207]	[-0.03, 0.847]
Time investment residual			0.231	-0.025
			[-0.093, 0.565]	[-0.346, 0.425]
Complementarity parameter	0.160	-0.041	0.067	-0.091
	[-0.083, 0.303]	[-0.23, 0.098]	[-0.11, 0.225]	[-0.235, 0.279]
Elasticity of substitution	1.191	0.961	1.071	0.917
	[0.924, 1.435]	[0.813, 1.109]	[0.901, 1.291]	[0.81, 1.388]
Over-identification test (p-value)			0.876	0.956
F-test two control functions (p-value)			0.147	0.209
F-test four control functions (p-value)				0.273

Note: Estimates based on 10,000 draws from estimated joint distribution of factors and instruments. Bootstrap 95% confidence intervals in square brackets based on 1,000 replications.

has a p-value of 0.27. Individually in the two production functions the p -values for the joint exogeneity test are 0.21 and 0.15 in the cognitive and socio-emotional production functions respectively.

When we assume investments to be exogenous, their impact on child development is very low, a fact that seems in contradiction with the importance that parents appear to give to such investments. When we allow for endogeneity, their impact is larger, albeit a lot more imprecise.

To investigate this further we experiment with more parsimonious, and possibly more precise, alternative specifications for the production function of cognitive skills and report their estimates in Table 5. In these experiments, we set the elasticity of substitution to be one (Cobb-Douglas) as implied by the estimates in Table 4. In the cognitive production function, we also eliminate time investments and keep only material investment - time investments seem to play a role only for socio-emotional development. Note that with just one investment in the cognitive skills production function the instruments are now strong and the weakness of the Cragg-Donald rank test is no longer an issue.

In the first two columns of Table 5, we present the estimates of a Cobb-Douglas specification without time investments. Whether or not we treat material investments as exogenous (column 1) or endogenous (column 2), the estimates of the production function are not very different from those in Table 4, except for the fact that in column 2 the treatment effect through TFP is smaller and insignificant. In these specifications, maternal cognitive skills only enter significantly when we assume investment is exogenous, and maternal socio-emotional skills never seem to matter. The p-value that both can be excluded from column 2 is 0.30. However, as we saw in Table 3, the maternal cognitive factor has a very strong influence on investments.

Based on this, we try two alternative specifications in which we treat investments as endogenous. First, we estimate a specification where we exclude the treatment indicator from the production function (column 3) and use it as an instrument. The instruments now are unquestionably strong and the impact of investment is much higher. Then in column 4 we present results where we exclude maternal skills, imposing that they only matter through the initial conditions of the child and through investments (where their impact is substantial).

Table 5: Alternative specifications for the production functions

Control function	Cognitive skill				Socio-emotional skill
	(1) No	(2) Yes	(3) Yes	(4) Yes	(5) No
TFP	0.004 [-0.011,0.011]	0.004 [-0.011,0.012]	0.027 [-0.008,0.075]	0.004 [-0.011,0.012]	-0.004 [-0.016,0.016]
TFP * Treatment	0.094 [0.017,0.191]	0.067 [-0.026,0.219]		0.031 [-0.059,0.154]	-0.022 [-0.106,0.07]
Baseline cognitive skill	0.638 [0.567,0.759]	0.622 [0.539,0.765]	0.61 [0.528,0.739]	0.602 [0.523,0.736]	0.098 [0.009,0.19]
Baseline socio-emotional skill	0.035 [-0.065,0.124]	0.039 [-0.057,0.127]	0.057 [-0.055,0.133]	0.056 [-0.041,0.155]	0.468 [0.372,0.617]
Maternal cognitive skill	0.219 [0.089,0.35]	0.135 [-0.118,0.427]	0.074 [-0.149,0.305]		0.014 [-0.152,0.151]
Maternal socio-emotional skill	0.047 [-0.009,0.124]	0.039 [-0.017,0.123]	0.035 [-0.024,0.12]		0.019 [-0.055,0.1]
Material investment	0.089 [0.025,0.149]	0.197 [-0.193,0.486]	0.276 [-0.001,0.516]	0.345 [0.208,0.466]	0.14 [0.06,0.235]
Time Investment					0.139 [0.038,0.233]
Number of children	0.038 [-0.025,0.098]	0.041 [-0.029,0.101]	0.045 [-0.022,0.105]	0.04 [-0.027,0.094]	0.135 [0.072,0.19]
Material investment residual		-0.115 [-0.428,0.275]	-0.194 [-0.466,0.097]	-0.263 [-0.418,-0.091]	
<i>Over-identification test</i>					
p-value		0.778	0.712	0.780	
<i>Goodness-of-fit: Gap in output between treated and control as</i>					
(a) Measured in the data	0.112	0.112	0.112	0.112	0.094
(b) Predicted by the model	0.112	0.112	0.066	0.112	0.094
(c) Predicted by the model but setting TFP * Treatment to 0	0.018	0.045	0.066	0.081	0.105

Estimates based on 10,000 draws from estimated joint distribution of factors and instruments. Bootstrap 95% confidence intervals in square brackets based on 1,000 replications.

In doing this we are able to keep the treatment indicator in and still have very strong instruments; however, its coefficient is less than half the size and insignificant (albeit more precisely estimated). Indeed if we remove it the results remain completely unchanged (not

shown) with some increase in precision. In this specification, exogeneity of investment is rejected, as indicated by the significance of the residual.³¹ Importantly, the tests of over-identifying restrictions has a p-value of 0.78, even with instruments that have a p-value of zero. Indeed if we add marital status in the production function it is never significant. This is also true if we add marital status to the specification shown in column 4; when included there, nothing changes and the marital dummy has a coefficient of -0.019 with a confidence interval of (-0.075,0.036).

We now turn to the fit of the model. When the full Cobb-Douglas specification is used with exogenous investments, the model fits the impact of the intervention perfectly (see the bottom of Table 5, which shows that the model predicts a 0.112 shift in the cognitive factor) but its estimates imply that the change in material investment account for only 16% of the effect of the intervention on child's cognitive skills. In that specification, the role of investment in the production function is small. However, when we tighten the specification by excluding mother's skills, we find that investment plays a much stronger role and the change in investment resulting from the intervention explains 72% of the intervention's impact on children's cognitive skills (even when the treatment indicator is included).

Statistically it is hard to distinguish between these alternatives, which is why we present them all. However, we believe that the specification in column 4 best reflects the cognitive skills production function in this data and context. It is important to stress that the specification does not ignore the genetic or very early contribution of parents, but does assume that it is all reflected in the initial conditions - a first order Markov type assumption, as well as through investments. The fact that this specification (that also includes the treatment indicator) allows for a stronger role for investment in the child also seems consistent with the facts that higher skill mothers invest more and children with higher cognition receive more investments. Finally, our interpretation of the negative (and significant) coefficient on the investment residual is that parents compensate for adverse events by increasing their in-

³¹p-value 0 in column 4 and p-value of 0.081 in column 3

vestment - hence the downward bias in the investment coefficient when we do not control for endogeneity. Moreover the improvement in cognition from our intervention operates almost entirely through a shift in investments.

We also tried the same type of analysis with the production function for socio-emotional skills. We found no evidence that investments are endogenous for that function even with our most parsimonious model. We thus keep to the specification with no control functions but we impose a substitution elasticity of one (Cobb-Douglas) as implied by the estimates in Table 4. These final results are in the last column of Table 5. As with the cognitive production function, maternal skills have no direct impact on socio-emotional development, again operating only through investments and the initial conditions of the child. Finally, as for cognitive skills, there is no evidence that there is any direct effect of the intervention: socio-emotional skills were improved by the increase in time and material investments induced by the intervention.

6.3 Implications of the estimates

The coefficients in Table 5 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 development of future cognitive (socio-emotional) skills. Second, we find evidence of cross-productivity: the current stock of cognitive skills fosters the development of future socio-emotional skills (although the effect is small), but the reverse does not seem to be the case. 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.

As we saw earlier, the number of children in the household (at baseline) reduces parental investments, especially time investments. However, in the production function, we find that the presence of siblings at baseline improves the socio-emotional development of our subject child (who is the youngest child in the family in most cases): the interaction of children in

the household seems to help develop important capabilities.

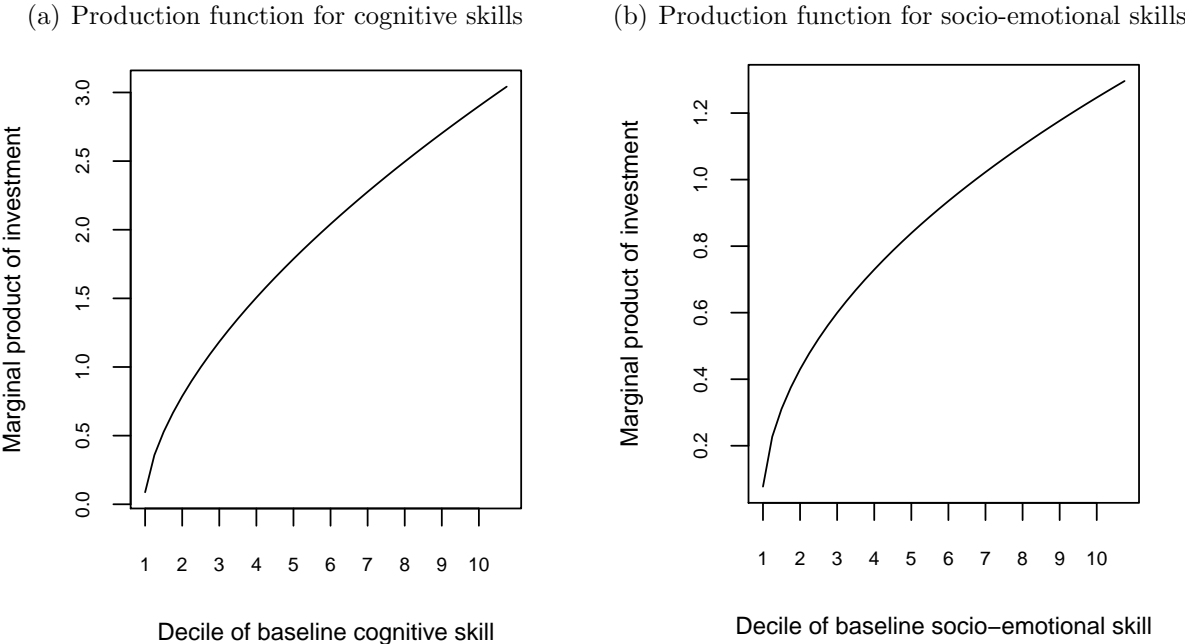
Finally, we turn to arguably the most important coefficients: the parameters associated with investments. The results show that material investments help develop both cognitive and socio-emotional skills; however time investments only influence the latter. The intervention increased the child’s cognitive factor by 0.112 log points and the socio-emotional one by 0.094 log points. Based on our preferred estimates the increase in parental investments induced by the intervention explains 72% of the improvement in cognition and all of the improvement in socio-emotional skills. Again it is important to associate these conclusions to the underlying measures: material investments relate mainly to toys and play materials, while time investments relate to time spent on stimulating activities with the child.

We have already established both directly through the experimental results and through the prism of our model that the intervention increased investments, as well as cognitive and socio-emotional skills. The key implication of the production function estimates is that improving investments can help child development substantially. These results show that the parenting component, where the home visitor directly involves the mother in developing the stimulation activities, is central to the intervention. While further experimentation is called for, the estimates suggest that home visits without the mother (or principal caregiver) present would not be as effective: in our model the direct effect of the intervention would show up as a change in the production function, i.e. as a new input. Change in investments reflects change in parental behavior only, since the underlying measurements only relate to what parents do outside the home visits.

The Cobb-Douglas specification implied by the data means that the inputs are complementary and provides further support to the concept of complementarities (Cunha et al., 2006; Heckman and Mosso, 2014). Given the metric we use for the latent factors, the return to investment is higher for children with better initial conditions. This is consistent with the fact that parents invest more in children with higher levels of cognitive development. To illustrate the extent of such complementarity, we plot in Figure 3 the marginal return

to material and time investments implied by the estimates of the model as a function of baseline cognitive and socio-emotional skills. Specifically, the graphs show how the effect of one standard deviation increase in material investment on cognitive skills (figure 2a) and how the effect of one standard deviation increase in both material and time investments on socio-emotional skills (figure 2b) varies by the initial level of these skills. The y-axis are in standard deviation units of the outcome.

Figure 3: Complementarity between investments and baseline skills



Notes: Figure 2.a (2.b) is based on the estimates of the production function for cognitive skills (socio-emotional skills) reported in Column 4 (5) of Table 5. The figures above are constructed by evaluating the increase in cognitive (socio-emotional skills) in standard deviation units resulting from an increase in one standard deviation of investments at different deciles of $\theta_{i,t}^C$ for (a) and $\theta_{i,t}^S$ for (b) and holding all remaining inputs of the production function at their mean values across the sample.

The fact that investments are more productive for children with higher levels of early cognition and socio-emotional skills also implies that the intervention had a lower effect on the more disadvantaged children, at least in the short-term, which raises the issue of how to better intervene on the most disadvantaged children and design interventions appropriate

for different ability levels.

While this finding may appear contradictory to the set of studies indicating that early interventions benefit low-achieving children the most (Bitler et al., 2014; Elango et al., 2016), one needs to allow for the differences in the populations concerned. Our intervention targets the 20% poorest children in Colombia. While these children do not live in extreme poverty, they may still be poorer and of lower ability at baseline than disadvantaged children targeted by programmes such as Head Start in the US. Our results imply that, in this subset of the population, those with a better start benefited more. However, one can imagine that with a population that extends more broadly in the socio-economic distribution, diminishing returns could set in unless perhaps we design an intervention better attuned to higher ability children.

Moreover, to our knowledge, very few studies are able to investigate the impact of interventions or of parental investments amongst children with different levels of baseline skills, and those that do actually find evidence of complementary effects (Aizer and Cunha, 2012). This suggests the possibility that, while interventions are indeed more effective among children who have stronger learning foundations, they achieve higher benefits among more disadvantaged children (and hence children with lower baseline skills) because they are able to compensate for the many risk factors that negatively affect their development.

7 Conclusion

Children from poor backgrounds accumulate developmental deficits from a very early age (Lancet, 2016; Rubio-Codina et al., 2015). Causes include not only the risky environments in which they live but also the lack of stimulation, which prevents the brain from developing its full potential. Such adverse early experiences are at the heart of the intergenerational transmission of poverty.

In this paper, we present results from an early childhood intervention carried out in Colombia that promoted suitable parenting and stimulation to children between one and

three years old. The intervention involved weekly home visits delivered by local women who had no prior knowledge of child development, but were trained to deliver a structured stimulation curriculum that progressed in difficulty. The evaluation by randomized controlled trial showed improvements in a number of developmental dimensions, including cognition and language. Importantly, it induced parents to invest more time and resources with their children.

We use data from the experiment to estimate a model of parental investments in time and resources and production functions for cognition and socio-emotional skills for their children. The aim of this model is to provide an interpretation of how the intervention affected child development and to improve our understanding of the development of child skills from a very early age. The model estimates trace some of the origins of social inequalities to the beginning of life: children with higher initial skills obtain more investments and, given their skills, mothers with higher levels of cognition invest more in their children. We show that the intervention increased both time and resource investments substantially.

Our best estimates of the production functions imply that these increased investments account for 72% of the intervention impact on cognition and 100% of its impact on socio-emotional skills. In other words, the program worked by inducing parents to invest more, rather than by making their investments more effective or through the direct impact of the home visits on the children.

Our study answers some important questions but raises many more, calling for further experimentation and analysis. For example, the results imply a complementarity between initial skills and the effectiveness of investments, suggesting that children in a better position at baseline benefited most. This emphasizes the need to understand how to better target and treat the most disadvantaged of the poor. Moreover, the analysis raises the question of how sustainable the effects of the intervention are and how salient improvements at this age are for longer term outcomes. This requires longer-term follow ups of the children participating in the intervention and calls for further research with systematic measurements

and interventions at various stages of life. Finally, a key lesson from our paper is the importance of parenting in promoting child development: interventions that improve parental inputs in poor environments can have substantial effects even with minimum resources, such as home made toys and good use of time.

[Click this link to obtain the Web Appendix.](#)

References

- Agostinelli, F. and Wiswall, M. (2016). Identification of dynamic latent factor models: The implications of re-normalization in a model of child development. Mimeo, Arizona State University.
- Aizer, A. and Cunha, F. (2012). The production of human capital: Endowments, investments and fertility. NBER Working Paper 18429.
- Almlund, M., Duckworth, A., Heckman, J. J., and Kautz, T. (2011). Personality, psychology and economics. In E. Hanushek, S. Machin, and L. Woessmann, eds., *Handbook of the Economics of Education*. Elsevier Science.
- 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.
- 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.
- Arcidiacono, P. and Jones, J. B. (2003). Finite mixture distributions, sequential likelihood and the em algorithm. *Econometrica* 71:933 – 946.
- Attanasio, O., Fernandez, C., Fitzsimons, E., Grantham-McGregor, S., Meghir, C., and Rubio-Codina, M. (2014). Using the infrastructure of a Conditional Cash Transfer programme to deliver a scalable integrated early child development programme in Colombia: A cluster randomised controlled trial. *British Medical Journal* 349:g5785.
- Attanasio, O., Fitzsimons, E., Gomez, A., Gutierrez, M.-I., Meghir, C., and Mesnard, A. (2010). Children’s schooling and work in the presence of a conditional cash transfer program in rural colombia. *Economic Development and Cultural Change* 181–210.
- 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.
- Bitler, M., Hoynes, H., and Domina, T. (2014). Experimental evidence on distributional impacts of head start. NBER Working Paper 20434.

- Black, M. M., Walker, S. P., Fernald, L. C. H., Andersen, C. T., DiGirolamo, A. M., Lu, C., McCoy, D. C., Fink, G., Shawar, Y. R., Shiffman, P. J., Devercelli, A. E., Wodon, Q. T., Vargas-Barn, E., and Grantham-McGregor, S. (2016). Early childhood development coming of age: science through the life course. *Lancet*. 6736 1–14.
- Caldwell, B. and Bradley, R. (2001). *HOME Inventory and Administration Manual*. University of Arkansas for Medical Sciences, Little Rock, AK, third edition edition.
- Campbell, F., Conti², G., Heckman, J. J., Moon³, S. H., Pinto, R., Pungello, E., and Pan¹, Y. (2014). Early childhood investments substantially boost adult health. *Science* 343:1478 – 1485. DOI: 10.1126/science.1248429.
- Campbell, F. A. and Ramey, C. T. (1994). Effects of early intervention on intellectual and academic achievement: A follow-up study of children from low-income families. *Child Development* 65:684 – 698. Doi:10.2307/1131410 Medline.
- 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.
- Cattell, R. (1966). The scree test for the number of factors. *Multivariate Behavioral Research* 1:629–637.
- 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.
- DelBoca, D., Flinn, C., and Wiswall, M. (2014). Household choices and child development. *Review of Economic Studies* 81:137–185.
- Duflo, E., Hanna, R., and Ryan, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review* 102(4):1241–78.
- Dunn, L., Padilla, E., Lugo, D., and Dunn, L. (1986). *Manual del Examinador para el Test de Vocabulario en Imgenes Peabody (Peabody Picture Vocabulary Test), Adaptacin Hispanoamericana (Hispanic-American Adaption)*. Pearson Assessments, Minneapolis, MS.
- Elango, S., García, J., Heckman, J., and Hojman, A. (2016). *Early Childhood Education*, volume Economics of Means-Tested Transfer Programs in the United States, volume 2. University of Chicago Press, Chicago, IL.

- Fernald, L., Kariger, P., Engle, P., and Raikes, A. (2009). *Examining child development in low-income countries: A toolkit for the assessment of children in the first five years of life*. The World Bank, Washington DC.
- Frongillo, E., Sywulka, S., and Kariger, P. (2003). UNICEF psychosocial care indicators project. Final report to UNICEF. Mimeo, Cornell University.
- Gelber, A. and Isen, A. (2010). Children’s schooling and parents’ behavior: Evidence from the head start impact study. *Journal of Public Economics* 101:25 ? 38.
- 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*.
- Gorsuch, R. (1983). *Factor Analysis*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Gorsuch, R. (2003). Factor analysis. In *Handbook of psychology: Research methods in psychology*, volume 2, 143–164. John Wiley & Sons, Inc., Hoboken, NJ.
- Grantham-McGregor, S., 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.
- 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.
- Hamadani, J., Tofail, F., Hilaly, A., Huda, S., Engel, P., and Grantham-McGregor, S. (2010). Use of Family Care Indicators and their relationship with child development in bangladesh. *Journal of Health Population and Nutrition* 28(1):23–33.
- 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. and Mosso, S. (2014). The economics of human development and social mobility. *Annual Reviews of Economics* 6:689–733.
- 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.
- Horn, J. (1965). A rationale and a test for the number of factors in factor analysis. *Psychometrika* 30:179–185.

- Hu, Y. and Schennach, S. (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica* .
- Jackson-Maldonado, D. (2011). Spanish cdi-iii development and preliminary results. Presentation at the International Congress for the Study of Child Language, Montreal.
- Jackson-Maldonado, D. and Conboy, B. (2011). Preliminary norms and validity for the Spanish CDI-III. American Speech-Language-Hearing Association (ASHA) Annual Conference. San Diego, CA.
- 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.
- Jackson-Maldonado, D., Thal, D., Marchman, V., Newton, T., Fenson, L., and Conboy, B. (2003). *MacArthur Inventarios del Desarrollo de Habilidades Comunicativas: User's guide and technical manual*. Brookes, Baltimore, MD.
- Jacoby, H. G. (2002). Is there an intrahousehold 'flypaper effect'? evidence from a school feeding programme. *The Economic Journal* 112:196 – 221.
- Kaiser, H. (1960). The application of electronic computers to factor analysis. *Educational and Psychological Measurement* 20:141–151.
- 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.
- Lancet (2016). Early child development in developing countries 2016 .
- Lu, C., Black, M. M., and Richter, L. M. (2016). Risk of poor development in young children in low-income and middle-income countries: an estimation and analysis at the global, regional, and country level. *Lancet Global Health* 4 e916?e922.
- Nielsen, E. R. (2015). Achievement gap estimates and deviations from cardinal comparability. Finance and Economics Discussion Series 2015-040. Washington: Board of Governors of the Federal Reserve System, <http://dx.doi.org/10.17016/FEDS.2015.040>.
- 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.
- Putnam, S. P., Jacobs, J., Garstein, M., and Rothbart, M. (2010). Development and assessment of short and very short forms of the Early Childhood Behavior Questionnaire.

- Radloff, L. (1977). The CES-D scale: A self-report depression scale for research in the general population. *Applied Psychological Measurement* 1(3):385–401.
- Raven, J. (1981). Manual for Raven’s progressive matrices and vocabulary scales. In *Research Supplement No.1: The 1979 British Standardisation of the Standard Progressive Matrices and Mill Hill Vocabulary Scales, Together With Comparative Data From Earlier Studies in the UK, US, Canada, Germany and Ireland*. Harcourt Assessment, San Antonio, TX.
- Romano, J. P. and Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. *Econometrica* 73:1237 – 1282.
- Rubio-Codina, M., Attanasio, O., Meghir, C., Varela, N., and Grantham-McGregor, S. (2015). The socioeconomic gradient of child development: Cross-sectional evidence from children 6-42 months in bogota. *Journal of Human Resources* 50(2):464–483.
- Schennach, S. (2007). Instrumental variable estimation of nonlinear errors-in-variables models. *Econometrica* 75:201–239.
- Shennach, S. (2004). Estimation of nonlinear models with measurement error. *Econometrica* 72:33–75.
- 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.
- Velicer, W. (1976). Determining the number of components from the matrix of partial correlations. *Psychometrika* 41:321–327.
- 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. P., Chang, S. M., Powell, C. A., and Grantham-McGregor, S. M. (2005). Effects of early childhood psychosocial stimulation and nutritional supplementation on cognition and education in growth-stunted Jamaican children: prospective cohort study. *The Lancet* 366:1804 – 1807.
- Walker, S. P., Wachs, T. D., Grantham-McGregor, S. M., Black, M., Nelson, C., and Huffman, S. (2011). Inequality in early childhood: risk and protective factors for early child development. *Lancet* 378(1325-38).