

Estimating the Production Function for Human Capital: Results from a Randomized Controlled Trial in Colombia[†]

By ORAZIO ATTANASIO, SARAH CATTAN, EMLA FITZSIMONS, COSTAS MEGHIR,
AND MARTA RUBIO-CODINA*

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 parents' 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, emphasizing the importance of parenting interventions at an early age. (JEL I24, I28, J13, J24, O15)

The first five years of life lay the basis for lifelong outcomes (Currie and Almond 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 nonstimulating home environments preventing them from reaching their full potential, particularly in developing countries (Grantham-McGregor et al. 2007; Black et al. 2016; Lu,

*Attanasio: Department of Economics, Yale University, Institute for Fiscal Studies, FAIR, and NBER (email: orazio.attanasio@yale.edu); Cattan: Institute for Fiscal Studies and IZA (email: sarah_c@ifs.org.uk); Fitzsimons: UCL Institute of Education and Institute for Fiscal Studies (email: e.fitzsimons@ucl.ac.uk); Meghir: Department of Economics, Yale University, and NBER, IZA, and Institute for Fiscal Studies (email: c.meghir@yale.edu); Rubio-Codina: Inter-American Development Bank (email: martarubio@iadb.org). Hilary Hoynes was the coeditor for this article. We thank four anonymous referees for detailed comments. We have benefited from extensive discussions with Don Andrews, Joseph Altonji, and James Heckman. We are grateful to Sally Grantham-McGregor, who participated in designing and implementing the experiment we describe here. 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. The authors gratefully acknowledge the support of the British Academy (grant reference pf140104), the NIH (grant reference R01 HD072120), the Economic and Social Research Council's Centre for the Microeconomic Analysis of Public Policy at the IFS (grant reference ES/M010147/1), the Cowles Foundation and the ISPS at Yale. Attanasio's time was partly funded by an European Research Council (ERC) Advanced Grant (Grant Agreement No. 695300). This project has also received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement 646917, 240910, and 249612). The authors thank the Economic and Social Research Council (grant RES-062-23-1548), Inter-American Development Bank, World Bank, and International Growth Center for funding the intervention and data collection. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. All errors are the responsibility of the authors. The views expressed in this publication are those of the authors and not necessarily those of the Inter-American Development Bank, their Board of Directors, or the countries they represent.

[†]Go to <https://doi.org/10.1257/aer.20150183> to visit the article page for additional materials and author disclosure statements.

Black, and Richter 2016). Thus, children from poor backgrounds accumulate developmental deficits from a very early age (Rubio-Codina et al. 2015, 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 help overcome these detrimental factors and have positive effects on children's development in both the short and long term. Examples include the Jamaica study (Grantham-McGregor et al. 1991, Walker et al. 2011, 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 and evaluated by a cluster randomized controlled trial (CRCT). The intervention was based on the Jamaican model of psychosocial stimulation via weekly home visits based on the curriculum now known as Reach-Up and Learn, and also offered micronutrient supplementation. However, unlike the Jamaican program, it was designed to be scalable by training local women involved in the implementation of a large welfare program to administer the weekly home visits.

Attanasio et al. (2014) shows that stimulation led to highly significant improvements in cognition and language development measured immediately following the end of the intervention,¹ and that micronutrient supplementation did not affect any outcome observed in the data. The paper also reported impacts on two raw measures of the home learning environment. However, it did not put these results together into a model of cognitive and socio-emotional skill production that would allow assessing the (joint) role of parental investments and of the intervention in shaping child development in the earliest of years.

Building on these results, the main aim of this paper is to understand *how* the stimulation component of the intervention led to improvements in child development, described by both cognitive and socio-emotional skills. 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, taken as endogenous, 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. In so doing we also provide some of the first estimated models of parental investments and human capital production functions at such an early age in a context of poverty in a middle-income country.

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

¹Using the Bayley Scales of Infant and Toddler Development III (Bayley 2006), cognition improved by 0.26 standard deviations (SD) (p -value 0.002) and receptive language by 0.22 SD (p -value 0.032). These reported p -values are adjusted for testing 12 hypotheses.

example, provides 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 CRCT and distinguish between material investments (e.g., books and toys around the home) and quality time investments (e.g., time spent by an adult in the household on education activities with the child).²

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 quality time investments, which we treat as endogenous. Within this framework, we quantify by how much changes in parental investments contributed to improving child outcomes in the treatment group. We also test whether the intervention shifted the production function or otherwise changed its parameters, 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 behavior as investments in their children. To our knowledge, our sample is one of the largest ever collected with this type of data in the literature evaluating stimulation programs. Even with such rich data however, 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 in Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010). To deal with the possible econometric endogeneity of investments, we use an instrumental variable strategy, adapted to the latent factor structure of the model.

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. In line with the existing literature, we find that a child's current stock of skills fosters the development of future skills, although we do find mean reversion.³ Second, parental investments and in particular our measure of material investments are an important determinant of future cognitive and socio-emotional skills. This becomes even more evident once we control for the endogeneity of such investments in line with results from Cunha, Heckman, and Schennach (2010) and Attanasio, Meghir, and Nix (2019) in very different contexts.⁴

²See Del Boca, Flinn, and Wiswall (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.

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

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 United States. 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

With respect to the mechanisms underlying the impacts, we find that the intervention significantly increased parental investments among treated families compared to nontreated ones. This increase is the dominant reason underlying the observed impacts. The direct effect of the intervention, instead, is both small and insignificant. 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.

Beyond revealing important aspects of the process of human capital accumulation and parental investments at a very early age, the importance of our results lies in two key findings. First, parents reinforce interventions by investing more, not less: there is crowding-in of material and time resources. Second, the intervention works because of the increase in parental investments and there is little evidence of a direct effect. If the intervention did not induce parents to invest more, it would have had no discernible effect. This, together with the mean reversion result that governs the longer-term impact of the intervention, emphasizes the key importance of improving parenting practices for the success and longer term sustainability of early childhood interventions.

Along with Heckman, Pinto, and Savelyev (2013) and a few other papers (Todd and Wolpin 2006; Attanasio, Meghir, and Santiago 2012; Duflo, Hanna, and Ryan 2012), 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 CRCT 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) which documents the channels through which the Perry Pre-School program produced gains in adult outcomes. But our focus and methodology are different: Heckman, Pinto, and Savelyev (2013) performs 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.

The focus of our intervention is also different. Unlike the Jamaican intervention, which targeted malnourished children, and the Perry Preschool 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

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.

Transfer (CCT) program *Familias en Acción*, which covers the poorest 20 percent 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 I provides some background on the intervention. Section II describes the data and the factor model approach we take to extract error-free measures of children's skills, parental skills, and investments. Section III discusses the short-term impacts of the intervention and some suggestive evidence of its underlying mechanisms. Section IV presents our theoretical framework and its empirical implementation. Section V presents the estimates of the model and discusses their implications for our understanding of the intervention. Section VI concludes.

I. 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, which targets the poorest 20 percent of households in the country. The intervention lasted 18 months, starting in early 2010. Online 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, 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.

We implement two key innovations vis-à-vis the Jamaican intervention with scalability and sustainability in mind. Indeed, our program was specifically designed to go beyond the earlier small scale and tightly supervised efficacy trials. 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 delivered by women drawn from the local community, with no specific prior professional early years experience.

To this end, home visitors were drawn from a network of local women, created by the administrative setup of the CCT program. Every 50–60 beneficiaries elect a representative who is in charge of organizing social activities and acts as mediator between them and the program administrators. These women, known as *Madre Líderes* (MLs), are themselves beneficiaries of the program. Given they are selected by their peers, one can deduce that they enjoy the trust of the community

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

and are probably more entrepreneurial and proactive than the average beneficiary. In terms of specific characteristics, they are on average about 10 years older (37) and have about one more year of education (8.5) than the subject mothers. Their score on a vocabulary test is slightly higher than that of the mothers.⁶ Finally, as mentioned in the introduction, another distinct feature of our intervention is that we targeted a more general poor population, namely the beneficiaries of the CCT program, which in Colombia is offered to the 20 percent poorest segment of the population, 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 across central Colombia. After first stratifying into three large regions, 32 municipalities in each were randomly assigned to one of four groups: (i) psychosocial stimulation, (ii) micronutrient supplementation, (iii) both, and (iv) control. In each municipality, 3 MLs were selected and a subset of 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 MLs are such trusted figures in their communities, compliance was high and the average number of home visits made was 63, which is 81 percent of those scheduled. The attrition rate between baseline and follow-up was around 10 percent across treatment arms, and the difference in loss among the groups was not statistically significant.⁷

As reported in Attanasio et al. (2014), there was no significant impact of micronutrient supplementation on any child developmental outcomes. In this paper, therefore, we focus on the psychosocial stimulation arm of the program and 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. On average, households had been part of the CCT 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 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 labor market outcomes (Gertler et al. 2014). This in itself

⁶To measure vocabulary, we use the Peabody Picture Vocabulary Test (PPVT). The scores of the MLs were 28.2 versus 26.9 for the subject mothers in our sample. The difference has a p -value of 0.061.

⁷As we explain in Section IIA, our data at baseline and at follow-up come from a household survey and from direct assessments administered to children in a community center. The attrition rate for the household survey was 6.9 percent. The attrition rate for the direct assessments was 10.7 percent.

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. The approach is aimed at promoting attention to task, perseverance, and self-esteem. Similarly, there is a strong focus on labeling 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 (i.e., linking concepts, understanding cause-and-effect relations). All of these skills are associated with improved school readiness, school attainment, and other outcomes linked to socioeconomic success in life.

II. Data and Measurement System

In this section, we describe the data we use, which were collected around the evaluation of the parenting intervention mentioned above. We then discuss an effective way of extracting the relevant information from such rich data with a measurement system that explicitly takes into account the relationship between relevant factors and available measures and the presence of measurement error.

A. Data

The main data we use in this paper come 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 center. At the end of the paper, we also briefly discuss results from a second follow-up, two years after the end of the intervention, although we do not use those data here.

The household surveys contain information on an extensive set of socio-economic and demographic characteristics, alongside a wealth of information around parenting, parental characteristics, and maternal skills, including mothers' years of education, verbal ability, IQ, depressive symptoms, and knowledge of child development. Among others, we administered to mothers the Raven's progressive matrices to test for IQ and the CES-D 10-item scale to assess depressive symptoms.

To measure children's developmental outcomes, we collected data based both on maternal reports and on direct assessment of the child. 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

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 the Infant Characteristics Questionnaire (ICQ); and the attentional focusing and inhibitory control scales of the short versions of the Early Children's Behavior Questionnaire (ECBQ). 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, trained psychologists administered the Bayley Scales of Infant and Toddler Development III (Bayley 2006) in community centers.⁸ 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 household surveys also contain detailed information on parental investments. We used a slightly modified version of the UNICEF Family Care Indicators (FCI) (Kariger et al. 2012) which is based both on interviewer observations and maternal reports of the home environment. Specifically, this instrument includes interviewer observations of the types and numbers of play materials around the home and maternal reports of the types and frequency of play activities performed by the primary caregiver or any other adult older than 15 with the child in the last 3 days.

Examples of play materials include toys designed for learning shapes, toys that induce physical movement, coloring books, and picture books. Examples of play activities include reading or looking at picture books together, telling stories, and labeling items in the home. Importantly therefore, the instrument affords the possibility of distinguishing between parental investments in "materials" and in "quality time," which are likely to have different costs and perhaps different (but possibly complementary) effects on child development.

As we want to assess the role of parental investment and distinguish it from the direct role that the intervention might play on child development, in measuring materials we instructed the data collectors to gather separate information about those materials (such as certain books and toys) that were directly provided by the intervention. When estimating the distribution of the factors measuring the two different types of parental investments, we use the measures that are not linked directly to the intervention.

Finally, in addition to the survey data collected around the evaluation of the intervention, we also use additional data sources to obtain information on municipality-level variables that we use as instruments. In particular, we use data on prices (of toys and food) and on maternal childhood exposure to violence. Online Appendix B provides details on all the measures of child development, maternal skills, and parental investments collected as part of the survey and describes the auxiliary data sources we use to construct our instruments.

Online Appendix Table A.1 reports the baseline characteristics of children, their mothers, and their households. At baseline, the children are on average aged 18 months. About 10 percent of them were born premature and 14 percent of them

⁸ See Jackson-Maldonado, Marchman, and Fernald (2012) for MacArthur-Bates inventories; Bates, Freeland, and Lounsbury (1979) for the ICQ; and Putnam, Gartstein, and Rothbart (2006) for ECBQ and Bayley (2006).

were stunted. On average, their mothers are 26 years old and have about 7.5 years of education, and two-thirds of them are either married or cohabitating. 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 CES-D scale items), none of these differences are significant at all when we allow for multiple hypothesis testing using the Romano and Wolf (2005) procedure.

B. 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:

$$(1) \quad \theta_{t+1} = f_{t+1}(\theta_t, I_{t+1}, P_t, X_t, \eta_{t+1}),$$

where θ_t and θ_{t+1} are vectors of the 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_t are maternal skills measured at baseline, X_t is a vector of baseline household characteristics, such as household composition, and η_{t+1} are random shocks to child development. 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 $f_{t+1}(\cdot)$, reflecting, for example, better use of parental inputs.⁹

As Cunha and Heckman (2008) explains, an important obstacle to estimating such a function is that the skills and investments are inherently unobservable. The various measures described in Section IIA 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 could lead to severely biased results, whether the model is linear or not. We thus follow the approach of Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010)¹⁰ 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 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 $t - 1$ and t . 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

⁹We 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 cognition (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.

¹⁰More broadly, this approach relates to the identification and estimation of nonlinear models with classical measurement error (Schennach 2004, 2007).

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 (Gorsuch 2003, 1983). 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 online Appendix C). Assuming each measure is additively separable in the (log) of the latent factor it proxies,¹¹ 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:

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

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

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

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 mean zero measurement error terms which are assumed independent of the latent factors and of each other.¹²

An assumption we have made in writing the system above is that the measurement system is invariant between treated and controls. 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 only. As we show in online Appendix E, none of our results are sensitive to whether we allow for the measurement treatment to be affected by the treatment.

Because the latent factors are unobserved, identification 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 $k = \{C, S\}$, where M and T refer to material and time investments, and C and S refer to cognitive and socio-emotional skills. When it comes to the child's skills, we normalize the factor loading on the same measures at baseline and follow-up.¹³ 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 between the treatment group's location and that of the control group (which is set to 0) is taken to be the average effect of the treatment.

¹¹The measurement equations are specified in terms of the log latent factors. This ensures that the factors themselves only take positive values as required by the model.

¹²The 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.

¹³For cognitive skills, we define the scale based on the Bayley cognitive score both at baseline and follow-up. For socio-emotional skills, we normalize the factor loadings on the item measuring difficulty in child's temperament in the ICQ.

With the assumptions and normalizations already made and based on the Kotlarski theorem and further extensions, Cunha, Heckman, and Schennach (2010) shows that both the distribution of measurement errors and the latent factor distribution are nonparametrically identified so long as we have at least three measures with nonzero factor loadings corresponding to each latent factor.¹⁴ While these assumptions are sufficient for identification, some of them could be relaxed as shown in Cunha, Heckman, and Schennach (2010).¹⁵ For instance, the same measure could be allowed to load on several factors, provided 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.

A question 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, Heckman, and Schennach (2010), which provides a theoretical treatment.¹⁶ In our paper, we normalize the cognitive factor in both periods to 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. For socio-emotional skills we also normalize to the same ICQ item (whether the child is difficult) in both periods. The lack of long-term longitudinal data prevents us from converting these units to future earnings or other adult outcomes of interest.

C. Specification of the Measurement System

To implement the measurement system above, we first perform an exploratory factor analysis, reported in online Appendix C, to identify in a preliminary step the relevant measures and their allocation to factors. We then allocate measures observed in the data to particular factors, as is shown in Table 1. The factor loading on the first measure is normalized to 1 and thus this measure 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

¹⁴See also Schennach (2004, 2007); Hu and Schennach (2008); Carneiro, Hansen, and Heckman (2003); Heckman, Pinto, and Savelyev (2013); and Cunha and Heckman (2008).

¹⁵See also Carneiro, Hansen, and Heckman (2003) and Cunha and Heckman (2008).

¹⁶Cunha, Heckman, and Schennach (2010) provides a general theoretical treatment of anchoring and 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. Nielsen (2015) discusses using ordinal test scores to measure achievement gaps, and Agostinelli and Wiswall (2016) discusses how rescaling in a multiperiod production function can lead to biases in the estimation of the substitution elasticity.

TABLE 1—MEASUREMENT SYSTEM

Latent factor	Measurement	Survey	Signal	
			Control	Treatment
Child's cognitive skills at FU (θ_{t+1}^C)	Bayley: Cognitive	FU	0.78	0.79
	Bayley: Receptive language	FU	0.75	0.76
	Bayley: Expressive language	FU	0.78	0.79
	Bayley: Fine motor	FU	0.59	0.61
	MacArthur: Words the child can say	FU	0.64	0.65
	MacArthur: Complex phrases the child can say	FU	0.52	0.54
Child's cognitive skills at BA (θ_t^C)	Bayley: Cognitive	BA	0.70	0.70
	Bayley: Receptive language	BA	0.73	0.72
	Bayley: Expressive language	BA	0.75	0.74
	Bayley: Fine motor	BA	0.60	0.59
	MacArthur: Words the child can say	BA	0.45	0.44
Child's socio-emotional skills at FU (θ_{t+1}^S)	ICQ: Difficult (-)	FU	0.74	0.71
	ICQ: Unsociable (-)	FU	0.33	0.30
	ICQ: Unstoppable (-)	FU	0.59	0.55
	ECBQ: Inhibitory control	FU	0.73	0.69
	ECBQ: Attentional focusing	FU	0.27	0.24
Child's socio-emotional skills at BA (θ_t^S)	ICQ: Difficult (-)	BA	0.68	0.71
	ICQ: Unsociable (-)	BA	0.28	0.31
	ICQ: Unadaptable (-)	BA	0.35	0.38
	ICQ: Unstoppable (-)	BA	0.22	0.25
Material investment at FU (I_{t+1}^T)	FCI: Number of types of play materials	FU	0.94	0.97
	FCI: Number of coloring and drawing books	FU	0.17	0.29
	FCI: Number of toys to learn movement	FU	0.61	0.76
	FCI: Number of toys to learn shapes	FU	0.69	0.82
	FCI: Number of shop-bought toys	FU	0.61	0.76
Time investment at FU (I_{t+1}^T)	FCI: Number of types of play activities in last 3 days	FU	0.87	0.93
	FCI: Number of times told a story to child in last 3 days	FU	0.66	0.81
	FCI: Number of times read to child in last 3 days	FU	0.73	0.85
	FCI: Number of times played with toys in the last 3 days	FU	0.55	0.72
	FCI: Number of times named things to child in last 3 days	FU	0.56	0.73
Mother's cognitive skills at BA (P_t^C)	Mother's years of education	FU	0.54	0.50
	Mother's Raven's score (IQ)	FU2	0.54	0.51
	Mother's vocabulary	FU2	0.65	0.62
	FCI: Number of books for adults in the home	BA	0.39	0.36
	FCI: Number of magazines and newspapers in the home	BA	0.20	0.19
Mother's socio-emotional skills at BA (P_t^S)	CESD: Did you feel depressed? (-)	BA	0.70	0.73
	CESD: Are you bothered by what usually don't? (-)	BA	0.42	0.45
	CESD: Did you have trouble keep mind on doing? (-)	BA	0.49	0.52
	CESD: Did you feel everything you did was an effort? (-)	BA	0.45	0.49
	CESD: Did you feel fearful? (-)	BA	0.47	0.51
	CESD: Was your sleep restless? (-)	BA	0.34	0.38
	CESD: Did you feel happy?	BA	0.46	0.50
	CESD: How often did you feel lonely in the last 7 days? (-)	BA	0.51	0.55
CESD: Did you feel you couldn't get going? (-)	BA	0.48	0.52	

Notes: 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.

children aged 3–7 and therefore was only administered at follow-up.¹⁷ However, in both rounds, we use the same measure to normalize the child’s baseline cognitive and socio-emotional skills.

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 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 (two years after the end of the intervention). In both cases we checked and the intervention had no 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. To do so, we use the estimation method described in Attanasio, Meghir, and Nix (2019), which approximates the joint distribution of the latent factors by a mixture of normals (as in Cunha, Heckman, and Schennach 2010) and the measurement error distribution by a normal distribution. We report estimates of the factor loadings and distribution of measurement errors in online 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, in the control group, 78 percent of the variance in the *Bayley: Cognitive* item is due to signal, whereas only 52 percent of the variance in the *MacArthur: Complex phrases* item is due to signal. Overall, most measures are far from having 100 percent of their variance accounted for by signal, which illustrates the usefulness of the latent factor approach in modeling human capital accumulation and parental investments:

¹⁷The 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.

without such an approach, one would risk to obtain severely attenuated coefficients, masking the importance of investments and background variables on child development.

III. Short-Term Impacts on Child Outcomes and Parental Investments

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. Impacts on the latter provide the basic input to perform the mediation analysis discussed below to uncover the mechanisms behind the observed impacts on children's outcomes. We focus on the impact of the psychosocial stimulation component of our intervention because there were no significant impact of micronutrient supplementation on any child developmental outcomes (Attanasio et al. 2014).¹⁸

A. Impacts on Child Development and the Home Environment

Each panel of Table 2 reports 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. Impacts on the Bayley outcomes and two of the FCI measures were previously reported in Attanasio et al. (2014) and are repeated here for completeness. In addition to the impact on each measure, we also report the impact on the mean of the corresponding log latent factor. The results in panel A imply an increase of 0.250 standard deviations (SD) in cognitive development and an increase of 0.175 SD in receptive language, assessed using the Bayley.¹⁹ The cognitive factor summarizing all these effects shows a substantial and significant increase of 11.5 percent (0.115 log points) amongst the treated group relative to the control group. Panel B also shows that the intervention led to an overall improvement in socio-emotional development (p -value < 0.05).

In Figure 1, we plot the estimated 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 tests are reported in the figure and have been derived using the bootstrap).²⁰ Panels A and B 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. Panels C and D 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

¹⁸If we explicitly control for the fact that one-half of 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 online Appendix Table D.1).

¹⁹These 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. We report the impacts on each of the three treatment arms in online Appendix Table D.1.

²⁰The estimation method used for this purpose is based on Attanasio, Meghir, and Nix (2019).

TABLE 2—TREATMENT IMPACTS ON RAW MEASURES AND LATENT FACTORS

	Treatment effect		
	Point estimate	SE	Sample size
<i>Panel A. Child's cognitive skills at follow-up</i>			
Bayley: Cognitive	0.250	(0.063)	1,264
Bayley: Receptive language	0.175	(0.063)	1,264
Bayley: Expressive language	0.032	(0.062)	1,263
Bayley: Fine motor	0.072	(0.060)	1,262
MacArthur: Words the child can say	0.092	(0.064)	1,322
MacArthur: Complex phrases the child can say	0.058	(0.055)	1,322
Cognitive factor	0.115	(0.051)	
<i>Panel B. Child's socio-emotional skills at follow-up</i>			
ICQ: Difficult (-)	0.074	(0.045)	1,326
ICQ: Unsociable (-)	0.041	(0.054)	1,326
ICQ: Unstoppable (-)	0.032	(0.054)	1,326
ECBQ: Inhibitory control	-0.003	(0.058)	1,323
ECBQ: Attentional focusing	0.069	(0.049)	1,323
Socio-emotional factor	0.087	(0.044)	
<i>Panel C. Material investments at follow-up</i>			
FCI: Number of types of play materials	0.215	(0.064)	1,326
FCI: Number of coloring and drawing books	-0.133	(0.056)	1,326
FCI: Number of toys to learn movement	-0.048	(0.065)	1,326
FCI: Number of toys to learn shapes	0.416	(0.088)	1,326
FCI: Number of shop-bought toys	0.024	(0.061)	1,326
Material investment factor	0.227	(0.069)	
<i>Panel D. Time investments at follow-up</i>			
FCI: Number of types of play activities in last 3 days	0.277	(0.050)	1,326
FCI: Number of times told a story to child in last 3 days	0.138	(0.060)	1,326
FCI: Number of times read to child in last 3 days	0.362	(0.062)	1,326
FCI: Number of times played with toys in last 3 days	0.175	(0.060)	1,326
FCI: Number of times named things to child in last 3 days	0.137	(0.048)	1,326
Time investment factor	0.302	(0.068)	

Notes: All scores have been internally standardized nonparametrically for age and are expressed in standard deviation units (see online Appendix B for details about the measures and the standardization procedure). Measures followed by (-) have been reversed so that a higher score refers to better behavior. The effects relating to the *latent factors* are in log points. Coefficients and standard errors clustered at the municipality level (in parentheses) are 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 the child's sex, tester effects, and baseline level of the outcome.

distribution. For the socio-emotional factor, however, the shift occurs mainly for children below the median.

B. Suggestive Evidence on the Mechanisms behind the Impacts

In panels E and F 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 densities between control and treated groups. Panels C and D of Table 2 focus on the mean impacts of the stimulation intervention on parental investments and indicate substantial impacts on several individual items of the FCI, as well as on the two latent factors measuring investments. Panel D of Table 2 shows that all types of time activities increase, but among play materials the increase is not uniform (panel C). Specifically, there is an increase in most toys but a reduction in coloring books,

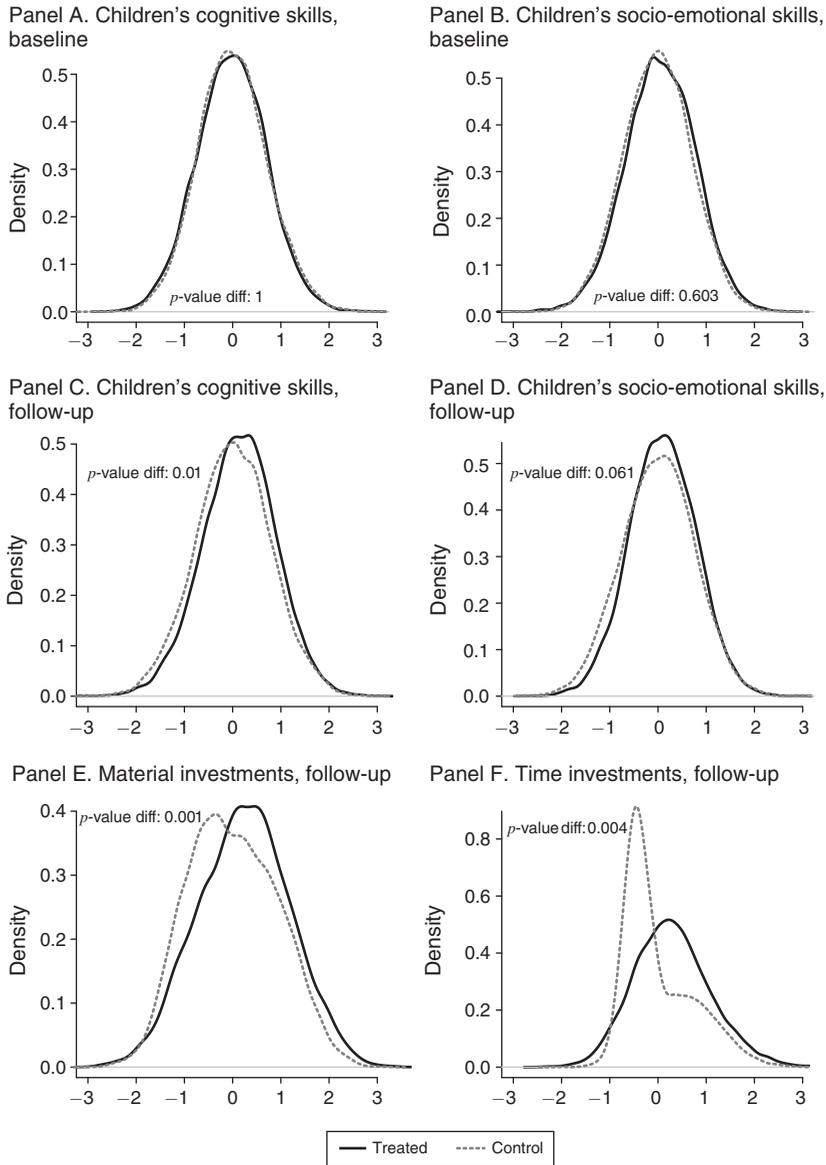


FIGURE 1. KERNEL DENSITIES OF LATENT FACTORS

Notes: 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.

which may reflect crowding out specially because the home visitors intentionally left picture books behind. The overall material investment factor registers a highly significant increase however.

As mentioned above, the measure of materials relates specifically to items provided by the parents. Although the home visitors were supposed to take away all intervention toys (with the exception of picture books) upon the completion of the last home visit, some were left behind at the end of the intervention as is evident

from the summary statistics reported in online Appendix Table D.3.²¹ However, since we are able to separately measure parental contributions from intervention materials, this does not pose any problem for our analysis, and whenever we refer to material investments we exclude intervention play materials and only keep items provided by the parents. The impact of the intervention materials, if any, will be captured by the treatment dummy; we return to this point when interpreting our results.

As we show above, the measures of parental inputs relating to materials and quality time both increased. As argued in Del Boca, Flinn, and Wiswall (2014), it may be important to distinguish these inputs because they can have different impacts on child development. On the other hand the inputs are clearly related and the two factors are correlated (0.62 in the control group). Thus, one could imagine a more parsimonious approach where they are combined into one investment factor. However, another good reason for keeping them separate relates to the way they are measured. Materials are actually observed and enumerated by the surveyors; quality time items are obtained by maternal self-reports, which may make them noisier measures of parental investment. This may also make them more likely to be subject to intervention bias due to the fact that mothers in the treatment group may exaggerate the extent to which they engage in developmental activities with the children. Without this implying that there is no information in the measures of time investments, these may be less reliable measures of parental behavior than material investments, which are directly observed and also likely to reflect actual effort by the parents. We will return to this important issue when interpreting the results.

IV. The Accumulation of Human Capital in the Early Years and the Role of the Intervention

To better understand the determinants of early childhood development and to explore the way that the intervention affected outcomes, we now specify a model of parental investments and child skill formation, where skills take two dimensions, namely cognitive and socio-emotional skills. We use such a model to inform the mediation analysis aimed at explaining the channels through which the intervention generated the impacts documented in Section III. An important element of the model is that parents can choose to invest in play materials and quality time.

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. However, rather than modeling investment choices resulting from the dynamic optimization of a household problem as in Del Boca, Flinn, and Wiswall (2014), we estimate a pair of reduced-form investment equations, which can be interpreted as an approximation to those derived (numerically) in a fully structural model. By not imposing all the restrictions from a specific

²¹The median number of days between the end of the intervention and data collection was 10. Almost all interviews were completed within 40 days. Very few households were interviewed a few days before the end of the intervention. Omitting these households from the analysis leaves the results completely unchanged.

structural model we do not have to take a stance on whether parents know the process of child development reflected in the structure of the production function.

The model we use allows us to characterize the process of early child development and provides a framework to understand the mechanisms that generated the intervention's impacts. One 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 total factor productivity (TFP) parameter 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 may not be captured by our measures and thus get embodied in the estimates of the production function parameters.²²

A second mechanism through which the intervention could generate impacts on child development 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 homemade 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 some form of investment either in itself (effectively an in-kind transfer). Such crowding-out of private resources is a standard concern in programs that target children.

Finally, the intervention could also have affected maternal cognitive or, more plausibly, socio-emotional skills. Many of the mothers (37 percent) were depressed at baseline according to the CES-D 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 the control and treatment groups 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.²³

A. *The Production Function for Human Capital*

We consider a production function of human capital that maps initial conditions, parental investment of different types and other factors on two different dimensions of child development. In particular, we assume that the stock of skills of child i 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

²²We 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 an adult in the household has read to the child in the last three 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.

²³The effect of the intervention on the principal component factor of the CES-D scale items at follow-up is 0.13 standard deviations (with a p -value of 0.12), which, given the way the factor is defined, is indicative of an improvement but too insignificant to rely upon.

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. As with skills, parental investments I_{it+1} can be a multi-dimensional vector. We denote material investments by I_{it+1}^M and time investments by I_{it+1}^T .

Following our own earlier experimentation, we assume the production function for each of the two skills is Cobb-Douglas, so we can write the technology of formation for skill k as follows:²⁴

$$(5) \ln(\theta_{idt+1}^k) = A_d^k + \gamma_{1d}^k \ln(\theta_{it}^C) + \gamma_{2d}^k \ln(\theta_{it}^S) + \gamma_{3d}^k \ln(P_{it}^C) + \gamma_{4d}^k \ln(P_{it}^S) \\ + \gamma_{5d}^k \ln(I_{idt+1}^M) + \gamma_{6d}^k \ln(I_{idt+1}^T) + \gamma_{7d}^k \ln(n_{it}) + \eta_{it+1}^k, \quad k \in \{C, S\},$$

where C and S stand for cognitive and socio-emotional skills respectively. The term n_{it} is the number of children in the household and allows 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.²⁵ It is possible, on the other hand, that the presence of siblings dilute attention and resources. As we discuss below, such an effect could be captured by the investment functions. The term A_d^k is a factor-neutral productivity parameter or TFP and depends on the treatment status of the child (d) to capture the potential direct effect of the home-visitor stimulation during her weekly visit. Finally, all the parameters are specific to a particular skill.

B. 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.²⁶ The number of children in the household may dilute both resources and time devoted to the subject child. We also include a vector of variables Z_{it} , which determine investments but do not enter the production function. We discuss them below. The investment equations we estimate are

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

As implied by the subscript d all coefficients could change with the treatment, a hypothesis we directly test. The effect of background variables on parental investment, given child initial conditions, is an important potential source of socioeconomic gradients in child development. Moreover, the extent to which investments increase with child initial abilities is a reflection of parental beliefs

²⁴Cunha, Heckman, and Schennach (2010) uses a CES, while Cunha and Heckman (2008) uses a log-linear specification.

²⁵Since our subject children are 12–24 months old at baseline these are almost always older children.

²⁶The measure includes the subject child so the minimum is 1.

about the heterogeneity of returns to such investments as well as parental taste for redistribution among children.

C. Estimation and Mediation Analysis

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

Standard mediation analysis, as in Heckman, Pinto, and Savelyev (2013), takes all inputs as exogenous. It then considers various possible channels through which the intervention could affect outcomes and tries to establish which of them can explain the observed impacts on the outcomes. In our case, this approach would correspond to estimating the production functions by ordinary least squares (OLS), allowing the intervention to affect outputs (cognition and socio-emotional skills) directly, as well as indirectly through its impact on investments. One could then decompose the overall effect of the intervention into a direct effect, which could be interpreted as an improvement in productivity, and an indirect effect mediated by the increase in investment. Such an approach, however, can lead to misleading conclusions if investments are endogenous: if, for instance, ignoring endogeneity leads to underestimate the impact of investment on outputs, then the channels through which the intervention can affect outcomes will be potentially misinterpreted. To deal with the endogeneity of investment we use instrumental variables. We discuss our choice of instruments in the next subsection.

To estimate the model we proceed in three steps. First, as mentioned earlier, we estimate the covariance structure and the factor loadings of the latent factors based on the covariance structure of the observed measures. This step requires no distributional assumptions on the latent factors or the measurement error, but only relies on the restrictions embedded in the measurement system and discussed earlier. For estimation however, we assume that the latent factors are distributed as a mixture of two normal distributions and that measurement error is normally distributed. In the second step, we use the estimates of the measurement system to predict Bartlett factor scores for each individual in the data. In the final and third step, we use these predicted scores as observable data to estimate the parental investment and production function equations.

The third step requires correcting for the measurement error introduced from the fact that we use predicted values of the latent factors instead of the actual ones. To do so, we adapt the correction method described in Heckman, Pinto, and

Savelyev (2013) in the context of instrumental variables. Using the predicted factor scores and instruments, we estimate the investment equations and the reduced form of the production functions, where we obtained the latter by substituting material and time investments in the production function with their relevant first-stage equations. We correct the investment and reduced-form coefficients using the method described in Heckman, Pinto, and Savelyev (2013) and finally recover the structural coefficients of the production function using a minimum distance estimator (Rothenberg 1971).

We compute 95 percent confidence intervals and critical values for test statistics using the cluster bootstrap, where the entire estimation procedure is replicated 1,000 times. We cluster at the municipality level which was the randomization unit in our experiment.

D. Choice of Instruments

When investments are treated as endogenous, identification requires instruments that are relevant and can be excluded from the production function, under reasonable assumptions. A potential instrument is the intervention itself, which was randomized. However, the fact that we wish to test whether the intervention affected the production function directly as one of the possible channels through which it operated precludes us from using it as an identifying variable that is excluded from the production function. Moreover, because we have two endogenous variables (material and time investments), we need at least two instruments so the randomization alone would not be enough to identify the model anyway.

Consistent with a standard model of parental investment, we expect material investments to be related to the prices of relevant goods. Specifically, we use as instruments the average log price of toys and the average log price of food items in the municipality of residence. We assume that the variability of prices across communities is unrelated to factors affecting the development of cognition and socio-emotional skills of children.

To provide corroborative evidence of the validity of these exclusion restrictions, in Table 3 we present regressions of the log price of food and the log price of toys on various baseline characteristics, which, reasonably, should not be affected by these prices. The coefficients are jointly insignificant as shown by the p -values at the bottom of the table. Moreover, the associated coefficients are very small and economically unimportant. Marriage is individually significant, but when all other variables are included in the model, it is not. In other versions of the model we also include the marriage indicator in the production function and none of the conclusions we draw are affected. Indeed, such a variable is completely insignificant in the production functions, further strengthening our conclusions.

Finding instruments for *quality time* is more challenging. This input may reflect more the way parents spend time with their children and the type of activities they engage in than the amount of time spent. This intuition is confirmed by results in Table 4, which show that the intervention had no impact on maternal labor supply: the impact of the treatment on both employment and weekly hours is small and insignificant, although it had large effects on time investments. This evidence

TABLE 3—BALANCE TEST FOR THE INSTRUMENTAL VARIABLES

	log toy price	log food price	Conflict
Mother's cognitive skill	0.021 (0.009)	0.007 (0.007)	-0.004 (0.003)
Mother's socio-emotional skill	-0.006 (0.008)	-0.014 (0.006)	0.005 (0.005)
Mother is married	0.013 (0.016)	0.030 (0.012)	0.017 (0.008)
Wealth index	-0.003 (0.007)	0.002 (0.006)	0.002 (0.004)
Terrorism	0.015 (0.014)	-0.012 (0.013)	0.005 (0.009)
Constant	8.025 (0.028)	8.062 (0.017)	0.049 (0.007)
Observations	1,010	1,023	1,023
R^2	0.011	0.021	0.017
F -statistic	1.656	2.147	1.438
F -test p -value	0.154	0.0671	0.219

Notes: All right-hand-side variables measured at baseline. Asymptotic standard errors in parentheses allowing for clustering at the municipality level. *Terrorism* is the number of terrorist attacks between conception of child and baseline. *Conflict* is the number of conflicts against civil population divided by the municipality population (in thousands) when the mother herself was a child.

suggests that the margin of adjustment is not between work and quality time with children, but rather between quality time versus other household production activities or leisure, which are excluded from our measures of time investments. Consequently, measures of the opportunity cost of time, such as village-level female or male wages, are unlikely to have much explanatory power.²⁷ Below we also show that quality time investment is also unrelated to the relative prices we use as instruments for material investment.

To instrument time investments, we instead focus on variables that have the potential to affect the willingness and ability of mothers to engage with their children. Specifically, we exploit the fact that Colombia has a long and well-documented history of civil conflict that has affected large parts of the country and, in particular, rural areas. It is well documented that exposure to violence can cause emotional detachment, which can impede or make subsequent interaction with one's own children harder (Betancourt 2015; Creech and Misca 2017, for related evidence). This leads us to consider maternal exposure to past conflict as a potential instrument for current quality time investment, where we exploit variation in the intensity of conflict across municipalities.²⁸

²⁷In earlier versions of the paper, we demonstrate that indeed village-level female or male wages do not predict time investments.

²⁸Specifically, our instrument is defined as the number of conflicts against the civil population divided by population (in thousands) in a given municipality when the mother was a child.

TABLE 4—IMPACT OF THE INTERVENTION ON MATERNAL LABOR SUPPLY

	Employment	Hours of work
Treatment dummy	−0.037 (0.085)	0.208 (1.793)
Demographic controls		
Mother's age	0.193 (0.044)	3.839 (0.802)
Mother's age squared	−0.003 (0.001)	−0.054 (0.013)
Years of education	0.056 (0.012)	0.921 (0.214)
Number of children = 2	−0.213 (0.098)	−4.429 (1.866)
Number of children = 3	−0.397 (0.124)	−5.905 (2.382)
Number of children = 4	−0.230 (0.166)	−2.057 (3.158)
Number of children = 5	−0.664 (0.236)	−13.711 (3.612)
Number of children ≥ 6	−0.346 (0.244)	−0.169 (5.107)
Constant	−3.457 (0.659)	−48.390 (11.913)
Observations	1,210	1,200

Notes: The table reports estimates of a regression of the measure of labor supply on the treatment dummy, mother's age, education, and number of children (all measured at follow-up), and interviewer fixed effects. *Employment* is a dummy taking the value 1 if the mother reports to work at follow-up and 0 if she does not; *Hours of work* measures her hours of work (0 if the mother does not work). Default number of children is 1, which is the subject child. Asymptotic standard errors in parentheses allowing for clustering at the municipality level.

As with our price instruments, our identifying assumption is that the variation in conflict across municipalities is orthogonal to unobservable factors affecting the development of the cohort of children we are studying. Although this assumption is not testable, in Table 3 we show that the incidence of conflict is not associated with mothers' baseline characteristics or with household wealth: all coefficients are very small and insignificant with a joint p -value of 0.22. Among these variables we also include the incidence of terrorist activities around the time the child was born: while the conflict itself was mainly over in the sampled communities by the time of our experiment, there were sporadic terrorist attacks during the period of the intervention. We included this variable to check that our measure of maternal childhood exposure to conflict is not related to current violence, which could have a direct impact on the child. The results in Table 3 strongly suggest that they are not. Thus, as is evident from these results, the past incidence of conflict is not associated with baseline characteristics relevant for child development. However, as we show below, mothers' childhood exposure to conflict is a strong determinant of the quality time that they spend with their child.

V. Results

We now report our empirical results. We start with estimates of the investment function before moving on to estimates of the production functions. The latter allow us to investigate what generates the impacts of the intervention on child development, a key question that influences the design of such programs. In what follows we report confidence intervals for all parameters. These have been computed using the block bootstrap accounting for the entire estimation procedure, and the cluster structure of our data. Whenever we present test statistics we compute the p -values using the bootstrap.

A. Estimates of the Investment Functions

The estimates of the investment equations are reported in Table 5. The first two columns of the table report estimates of the material and time investment equations, where we use toy price, food price, and maternal childhood exposure to conflict as exclusion restrictions. The third column reports estimates of the material investment equation, where we only include prices as exclusion restrictions. This first stage will be used to estimate a production function for cognitive skills that does not include time investments as an input (and that will turn out to be our preferred specification).

The first striking result is the impact of treatment on investments: it increases materials by 21 percent and time by 32 percent, and both effects are highly significant. The results reported in Table 5 exclude interactions of the treatment parameter with the remaining variables, which were found to be insignificant.²⁹ These estimates of the impact of the intervention on investments are driven by the experimental design and do not require any of the assumptions necessary for the identification of the production functions.

The fact that the intervention increased the quality time and resources that parents provide to children is important because it shows that parents are willing to reinforce the intervention. While we already showed some evidence of crowding out for individual items, overall the opposite seems to be happening. From a policy perspective this is a major conclusion that should encourage further interventions in such contexts. As we shall see below, this increase in parental investment is the key source of success of the intervention.

Turning now to the other regressors, we find that parents invest more resources in children with a higher baseline level of cognition (elasticity of 0.13) but the child's baseline socio-emotional skills have no impact on either type of investment. The elasticity of both material and time investments with respect to maternal cognition is very high and particularly so for the former; however, the mother's socio-emotional skills only affect material investments significantly and the impact is very small. The number of other children in the household at base-

²⁹The estimates where all parameters of the investment functions are allowed to vary with treatment are shown in online Appendix Table E.2. 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 0.369 and the p -value for the time investment equation is 0.099.

TABLE 5—ESTIMATES OF THE MATERIAL AND TIME INVESTMENT EQUATIONS

	Instruments: prices and conflict		Instruments: prices only
	Material investment	Time investment	Material investment
Intercept	−0.015 [−0.114, 0.078]	0.001 [−0.089, 0.089]	−0.013 [−0.11, 0.078]
Treatment	0.209 [0.038, 0.365]	0.318 [0.155, 0.48]	0.204 [0.037, 0.362]
log child's cognitive skill (t)	0.130 [0.016, 0.246]	0.068 [−0.044, 0.18]	0.132 [0.017, 0.25]
log child's socio-emotional skill (t)	−0.028 [−0.133, 0.087]	0.027 [−0.083, 0.145]	−0.030 [−0.131, 0.088]
log mother's cognitive skill	0.748 [0.582, 0.939]	0.349 [0.162, 0.498]	0.750 [0.583, 0.943]
log mother's socio-emotional skill	0.069 [−0.008, 0.139]	0.022 [−0.06, 0.108]	0.068 [−0.008, 0.139]
log number of children	−0.129 [−0.18, −0.077]	−0.128 [−0.186, −0.072]	−0.128 [−0.18, −0.078]
log toy price	−0.096 [−0.168, −0.027]	−0.020 [−0.085, 0.037]	−0.094 [−0.163, −0.026]
log food price	0.091 [0.006, 0.178]	0.042 [−0.026, 0.121]	0.091 [0.006, 0.178]
Maternal childhood exposure to conflict	−0.009 [−0.08, 0.063]	−0.089 [−0.139, −0.032]	
Rank test (p -value)		0.011	
Cragg-Donald test (p -value)		0.020	
Test of joint significance: F -statistic (p -value)			
Toy price, food price, conflict	10.47 (0.028)	13.42 (0.008)	
Toy price, food price, conflict, treatment	22.41 (0.001)	26.42 (0.001)	
Toy price, food price	10.47 (0.013)	1.53 (0.455)	10.57 (0.010)
Toy price, food price, treatment	22.40 (0.001)	20.56 (0.000)	22.55 (0.001)
Conflict, treat	5.94 (0.047)	21.80 (0.000)	

Notes: Dependent variables are the log of material and time investments at follow-up ($t + 1$). *Maternal childhood exposure to conflict* is the number of conflicts against civil population divided by the municipality population (in thousands) when the mother herself was a child. t refers to baseline/pre-treatment measurement. Ninety-five percent confidence intervals in square brackets. These as well as the p -values for the rank tests and all other tests are based on 1,000 bootstrap replications of the entire estimation process taking into account of clustering at the municipality level. *Rank test* is a test of the null hypothesis that the smallest eigenvalue of the 2×2 matrix $\beta'\beta$ is zero, where β is the 3×2 matrix of coefficients on *log toy price*, *log food price*, and *maternal childhood exposure to conflict* in the material and time investment equations. We present this alongside the Cragg-Donald test because in this context it is not clear which is the more powerful.

line reduce significantly both time and material investments: the elasticity of investment with respect to children is about -0.13 , which is consistent with a quantity/quality trade-off among children. Moreover, the results are in line with a model where parents choose investments taking into account complementarity with child cognitive skills.

Importance of Instruments.—The next set of variables reported in Table 5 explain investments and act as excluded instruments when we estimate the production functions and treat investments as endogenous. These are the prices of toys and

food and the level of conflict in the municipality when the mother was a child. The rank test we implement has a p -value of 0.011.³⁰ As an alternative, we also present the Cragg-Donald form of this test, which has a p -value of 0.020. These establish the strength of the instruments, allowing for the fact that there are two endogenous variables. As we will elaborate in Section VD, our Monte Carlo simulations show that these instruments are strong enough to avoid weak instrument bias.

The prices we consider are expected to affect material investment through the household budget constraint. Indeed, the p -value of a test of joint significance of these instruments is 0.013 in the material investment equation. As we would expect, toy price has a negative and significant impact on material investments. Food price, on the other hand, has a significant positive effect on material investment, implying that play materials and food are substitutes.

Understanding the determinants of quality time is harder, as argued above. Prices have no explanatory power: their joint p -value in the quality time investment equation is 0.46. In results not reported here, we also found that wages did not predict time investments either, supporting our argument that the opportunity cost of quality time is not the time spent in the labor market and that spending quality time with children does not necessarily require monetary resources, but rather perhaps knowledge of child development and a certain willingness to engage with the child. Maternal childhood exposure to conflict, on the other hand, has a strong and significant negative impact on quality time activities (p -value 0), although it has no impact on material investments. The latter result reinforces the idea that this variable is not merely picking up some omitted background factor.

Jointly, these instruments are strong enough to allow us to control for the endogeneity of investments in the production function. In some specifications, we also exclude treatment status from the production functions, thus implicitly using the treatment as an instrument. In this case, the instruments only become stronger.

Finally, in the last column of Table 5 we present a specification for material investment that excludes the conflict variable; the coefficients are almost identical to those in column 1, and the prices are jointly even more significant. We use this investment equation to estimate a model where only material investments enter the production function for cognitive skills.

B. The Production Function for Cognitive Skills

In Table 6, we report estimates of a Cobb-Douglas production function for cognitive skills. The first column presents results where investments are assumed to be conditionally exogenous; in the remaining columns investments are taken as endogenous.

The production function demonstrates a high level of persistence for cognition; however, socio-emotional skills do not affect cognition at this early age, and this result remains unchanged whether we treat investments as endogenous. Mothers' cognition and socio-emotional skills have a strong positive effect on cognitive development when we use OLS, but in the remaining columns, when we use IV for the investments, these effects disappear, implying that mothers' skills operate

³⁰See Robin and Smith (2000) and Blundell, Duncan, and Meghir (1998).

TABLE 6—ESTIMATES OF THE PRODUCTION FUNCTION FOR COGNITIVE SKILLS

Instruments:	OLS	IV			
	(1)	Prices, conflict (2)	Prices, conflict, treatment (3)	Prices (4)	Prices and treatment (5)
Intercept	−0.018 [−0.094, 0.053]	−0.019 [−0.111, 0.079]	0.007 [−0.089, 0.104]	0.003 [−0.091, 0.072]	−0.009 [−0.083, 0.058]
Treatment	0.083 [−0.025, 0.192]	0.049 [−0.12, 0.391]		−0.028 [−0.186, 0.156]	
log child's cognitive skill (t)	0.675 [0.589, 0.77]	0.648 [0.544, 0.795]	0.638 [0.522, 0.747]	0.626 [0.525, 0.746]	0.631 [0.533, 0.739]
log child's socio-emotional skill (t)	0.001 [−0.091, 0.087]	0.012 [−0.098, 0.143]	0.019 [−0.094, 0.14]	0.015 [−0.084, 0.126]	0.02 [−0.079, 0.127]
log mother's cognitive skill	0.213 [0.089, 0.35]	−0.075 [−0.456, 0.5]	−0.173 [−0.538, 0.201]	−0.102 [−0.495, 0.291]	−0.094 [−0.45, 0.21]
log mother's socio-emotional skill	0.103 [0.031, 0.173]	0.084 [−0.03, 0.163]	0.063 [−0.035, 0.151]	0.06 [−0.019, 0.152]	0.074 [−0.026, 0.155]
log number of children	0.042 [−0.01, 0.092]	0.085 [−0.07, 0.154]	0.084 [0.011, 0.163]	0.089 [0.002, 0.176]	0.086 [0.023, 0.164]
log material investment	0.088 [0.016, 0.157]	0.594 [0.025, 1.179]	0.784 [0.204, 1.383]	0.542 [0.041, 0.996]	0.516 [0.195, 0.946]
log time investment	0.038 [−0.051, 0.129]	−0.171 [−1.198, 0.312]	−0.311 [−0.985, 0.217]		
Goodness-of-fit: Gap in output between treated and control					
Measured in the data	0.115	0.115	0.115	0.115	0.115
Predicted by the model	0.114	0.132	0.084	0.095	0.117

Notes: Dependent variable is the log cognition of the child at follow-up ($t + 1$). t refers to baseline/pre-treatment measurement. At the top of each of the IV columns we state the exclusion restrictions used. Ninety-five percent confidence intervals in brackets based on 1,000 bootstrap replications of the entire estimation process taking into account clustering at the municipality level.

through the initial conditions of the child and through the investment decisions only. Finally, the number of other children in the family improves child cognitive development. This is particularly interesting because the investment equations show that additional children reduce both material and quality time investments. However the presence of other (for the most part older) children has a direct impact on child development: the elasticity of cognition with respect to children is 0.04–0.09 depending on the specification.

We now turn to the estimates of the direct treatment effect and the coefficients on parental investments, as measured by our “material investment” and our “quality time investment” factors. When we treat investments as exogenous and use OLS, we find that the impact of material investments on cognitive development is significantly different from zero, but that of time is insignificant. The direct treatment effect is large but very imprecisely estimated, to the extent that it is insignificant. The estimate of the impact of material investments increases dramatically when we treat investment as endogenous. Going from column 1 to column 2, this coefficient increases from 0.088 to 0.594; and although it is now estimated less precisely it remains significant at the 5 percent level. The change in the size of the coefficient points to parents compensating for negative shocks affecting their children.

The coefficient on time investment, instead, stays insignificant and its point estimate actually turns negative. The estimate of the direct effect of the intervention stays insignificant, although, as we discussed above it is strongly significant in the investment equations. For this reason, in column 3, we force the direct treatment effect to be zero. Despite the now unequivocal strength of the instruments (since we are also using the treatment allocation as an exclusion restriction in column 3), the results do not change: material investments enter with a large and significant coefficient (0.784) and time investments remain completely insignificant.³¹

Using Material Investments Alone.—Given the measurement issues for quality time, which is based on self-reports and discussed further below,³² and in light of the results presented thus far, in the next two columns we exclude time investments from the production function and we switch to using *only* prices as instruments so the relevant first stage investment equation is in the third column of Table 5. In one specification we include the treatment dummy reflecting the direct effect of the intervention (column 4) and in the other we exclude it (column 5). The results are essentially the same, though with improved precision, and present a clear message: material investments have a strong and positive impact, while the direct effect of treatment is small and completely insignificant.

Interpreting the Impact of the Intervention.—Based on these estimates, we now consider how the intervention affects outcomes through the lens of the production function. The possible channels include changes in the production function, a direct effect, and changes in parental investment inputs. Changes in the production function could happen for a number of reasons. First, the weekly session of the home visitor with the child, as well as any materials left behind, 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.

From the coefficients reported in column 1 of Table 6, which treats investment as exogenous, together with the increase in investment documented in Table 5, we see that increased investment accounts for about 25 percent of the impact of the intervention. The point estimate of the coefficient on treatment, instead, indicates that the intervention directly increased cognition considerably, by about 8.3 percent, although it is not precisely estimated. Therefore, according to this specification, parental investments play some role in mediating the intervention but there is a large direct effect, at least in terms of the point estimate.

A key problem with these results, though, is that they assume investments are invariant to unobserved shocks to child development that occur between baseline and

³¹In an earlier version of the paper we also used an indicator for whether the mother is married at baseline as an instrument. Following referees' comments, we no longer use it. It is worth mentioning though that while being married had a significant effect in the investment equations it was not at all significant when also included in the production functions. One interpretation is that married couples invest more in their children and this is the channel through which outcomes are affected in our context.

³²See also Section IIIB.

the end of the intervention 18 months later. The IV results presented in columns 2 to 5 of Table 6 address this issue and present a different story: the increase in material investments now explains a large fraction of the observed impact. Thus, a good description of cognitive development among children aged 3–4 is that it is driven by initial cognition (measured at ages 1–2), stimulation provided by older siblings, and material investments provided by parents, and the main channel through which the intervention affects cognitive development is by shifting these investments. The model in column 5 with material investments alone achieves a good fit of the data and captures all of the overall impact of the intervention. Finally, we find no evidence that any other coefficient of the production function changed as a result of the intervention. In online Appendix Table E.3 we allow all coefficients to vary by treatment status. The overall p -value, whether or not we include the interaction with the intercept, is 0.28.

These results highlight the importance of material investments but leave no role for quality time investments. This is perhaps surprising because one would expect quality time to be an important input as well. However, this result could at least partly reflect the measurement issue we discussed earlier (see Section IIIB) and that needs to be addressed in future research: material investments are measured through interviewer observations, while time investments are self-reported and hence carry the risk to be over-reported, particularly in the intervention communities where the importance of quality time has been continuously emphasized throughout the duration of the intervention. Material investments provided by the parents require real effort through making or buying toys, and they are likely to correlate strongly with actual activities carried out with the child. For example, providing a toy, whether it is homemade or bought, will likely include spending time with the child in this activity. The actual observation of play materials may therefore better reveal the developmental activities carried out by parents. On the other hand the materials left behind by the intervention and captured by the intervention/treatment dummy may not reflect actual engagement by the parent. This again emphasizes that the program works to the extent that it shifts actual parental behavior.³³

C. The Production Function for Socio-Emotional Skills

In Table 7, we present estimates of several specifications of the production function for socio-emotional skills. As with cognitive skills, we observe that the accumulation process of socio-emotional skills exhibits a substantial amount of persistence (regardless of the specification considered). The point estimates, however, are lower than in the case of cognition: about 0.50, compared to about 0.70 for cognitive skills. The intervention also has no direct effect in any specification. The lagged value of cognitive development is marginally significant in the various specifications, indicating a role for cognitive development in generating socio-emotional skills.

³³It may be interesting to consider whether similar impacts could be achieved by subsidizing prices of toys. However, the treatment effect on investment is clearly very large compared to the price effects. In a sense this reflects the fact that the intervention induces parents to change their behavior toward their children in quite a radical way that cannot reasonably be achieved by shifting prices.

TABLE 7—ESTIMATES OF THE PRODUCTION FUNCTION FOR SOCIO-EMOTIONAL SKILLS

Instruments:	OLS	IV			
	(1)	Prices, conflict (2)	Prices, conflict (3)	Prices, conflict, treatment (4)	Prices, conflict, treatment (5)
Intercept	−0.009 [−0.08, 0.063]	−0.02 [−0.087, 0.064]	−0.006 [−0.084, 0.066]	−0.02 [−0.096, 0.058]	−0.022 [−0.089, 0.051]
Treatment	−0.011 [−0.124, 0.093]	−0.071 [−0.318, 0.166]	−0.116 [−0.31, 0.115]		
log child's cognitive skill (t)	0.106 [0.018, 0.192]	0.074 [−0.019, 0.21]	0.09 [−0.017, 0.194]	0.094 [0.003, 0.22]	0.099 [0.002, 0.193]
log child's socio-emotional skill (t)	0.522 [0.403, 0.659]	0.513 [0.374, 0.672]	0.499 [0.389, 0.663]	0.516 [0.387, 0.669]	0.512 [0.403, 0.656]
log mother's cognitive skill	−0.077 [−0.217, 0.049]	−0.146 [−0.443, 0.297]	−0.15 [−0.349, 0.084]	−0.018 [−0.315, 0.328]	−0.083 [−0.231, 0.078]
log mother's socio-emotional skill	0.037 [−0.058, 0.119]	0.035 [−0.062, 0.141]	0.048 [−0.053, 0.135]	0.049 [−0.054, 0.134]	0.043 [−0.042, 0.126]
log of number of children	0.099 [0.047, 0.153]	0.127 [0.017, 0.236]	0.133 [0.034, 0.223]	0.101 [0.028, 0.168]	0.101 [0.036, 0.163]
log material investment	0.154 [0.06, 0.256]	0.015 [−0.621, 0.428]		−0.144 [−0.688, 0.34]	
log time investment	0.109 [−0.006, 0.213]	0.487 [−0.177, 1.258]	0.549 [−0.022, 1.147]	0.448 [−0.125, 1.057]	0.324 [0.025, 0.691]
Goodness-of-fit: Gap in output between treated and control Measured in the data	0.087	0.087	0.087	0.087	0.087
Predicted by the model	0.057	0.080	0.050	0.103	0.098

Notes: Dependent variable is the log socio-emotional skills of the child at follow up ($t + 1$). t refers to baseline/pre-treatment measurement. At the top of each of the IV columns we state the exclusion restrictions used. Ninety-five percent confidence intervals in brackets based on 1,000 bootstrap replications of the entire estimation process-clustering at the municipality level.

Parental background variables, such as mothers' cognitive and socio-emotional skills, are not significant predictors of children's socio-emotional skills. Instead, the number of siblings has a positive and significant impact on these skills, consistent with what we found for cognitive development. This result is robust across specifications.

As with cognitive skills, we experiment with a number of specifications. In none of the specifications does the direct treatment effect play any role: the coefficient is always negative and very imprecise. When we use OLS we find that both material and quality time investments enhance socio-emotional development. When we turn to IV (columns 2 to 5), none of the investment coefficients remain significant, with the exception of the results in column 5. There, we exclude material investments, which always enter with a negative coefficient when instrumented, and the treatment dummy, which never plays a role. In this case the results suggest that our measure of quality time can explain the improvement in socio-emotional skills resulting from the intervention.

Based on the OLS and using a specification that excludes the direct impact of the intervention (which is negative and insignificant) all the reported coefficients in column 1 remain unchanged. That specification explains 66 percent of the impact (it predicts a 0.057 log point improvement). The alternative specification in column 5 over-predicts the impact. We also tested whether the production function

differed beyond an intercept shift between treatment and control: the differences were completely insignificant, even for the more precise OLS specification with a p -value of 0.38. Whether we take the OLS results or the IV results presented in column 5, their implications are clear and similar: the intervention acted by improving parental engagement with their children but there is no evidence that it also had a direct impact on their development.³⁴

D. Are the Instruments Weak? Monte Carlo Evidence

Before discussing further the implications of the estimates, we first report results demonstrating that the instruments we use are strong enough to avoid biases due to a weak instrument problem. Weak instruments can lead to estimates that are severely biased toward OLS.³⁵ Although the F -statistics and the rank tests reported in Table 5 suggest that the instruments are sufficiently strong, these criteria have not been derived and validated for our nonstandard framework, which includes a first stage with factor analysis, two endogenous variables on the right-hand side and a cluster structure with intracluster spatial correlation. Instead of relying on unvalidated diagnostics, we investigate directly whether we have a weak instrument problem by performing a Monte Carlo simulation.

Using our parameter estimates as true values³⁶ we simulate investments and cognition. The simulations use the actual data on the exogenous variables,³⁷ which are kept fixed in repeated samples. We draw errors based on a covariance matrix that replicates the stochastic structure in the data, including the cluster structure. We use two alternative covariance matrices (Covariance A and Covariance B), the latter being particularly conservative.³⁸ By setting the coefficients of the first stage in the data generating process to be the same as the ones we estimate in the investment equations and using the appropriate covariance structure for the errors, we replicate data which leads to a first stage effectively equal in strength as the actual one in the data. All details are provided in online Appendix Section F.

We compare the assumed true values to the average estimates obtained in 1,000 simulated datasets. The difference is the bias when our estimator is applied to data of our structure and size. The results for some key parameters are shown in Table 8. Panel A compares the average estimates of the coefficients on the instruments for the first stage to the assumed true values used in generating the data. As expected they are estimated with no substantial bias, the average being very close to the true values.

Panel B of Table 8 shows the coefficients on the two investments in the cognitive production function. Under the column *OLS* we erroneously ignore the endogeneity

³⁴The experiment was stratified across three regions: Central, Oriental, and Cafetera. The probability of allocation to treatment was equal in each stratum and hence we do not need to control for stratum fixed effects in the experimental analysis. However, they may be relevant in the estimation of the investment equations and the production functions. Including the stratum fixed effects, as expected, weakens a bit the instruments, whose variation is geographic, although prices and conflict are still significant. However, the overall conclusions do not change. We present these additional results in online Appendix Tables E.6 and E.7.

³⁵See Bound, Jaeger, and Baker (1997); Staiger and Stock (1997); Andrews and Stock (2005) amongst others.

³⁶We interpret the coefficient of time investments in the cognitive production function as 0. In other nonreported simulations, we also allowed for a positive coefficient on time investment and none of the conclusions were affected.

³⁷Prices, maternal exposure to conflict, treatment status, maternal cognitive skill, and child's baseline cognitive skill.

³⁸Both covariance matrices are reported in online Appendix Table F.1.

TABLE 8—MONTE CARLO SIMULATIONS

Panel A. Coefficients on instruments in material and time investment functions

	Estimates on simulated data					
	True values		Covariance A		Covariance B	
	Material	Time	Material	Time	Material	Time
Toy price	-0.96	-0.02	-0.091 [-0.147, -0.036]	-0.019 [-0.079, 0.039]	-0.083 [-0.158, -0.009]	-0.02 [-0.113, 0.064]
Food price	0.91	0.042	0.089 [0.024, 0.154]	0.042 [-0.025, 0.108]	0.081 [-0.008, 0.167]	0.038 [-0.061, 0.133]
Conflict	-0.009	-0.089	-0.01 [-0.067, 0.042]	-0.094 [-0.156, -0.035]	-0.008 [-0.085, 0.068]	-0.084 [-0.17, 0.001]
Rank test (<i>p</i> -value)			0.042		0.029	
Cragg-Donald test (<i>p</i> -value)			0.06		0.021	

Panel B. Coefficients on material and time investments in the production function for cognitive skill

	Estimates on simulated data				
	True values	Covariance A		Covariance B	
		OLS	IV	OLS	IV
Material investment	0.594	0.11 [0.044, 0.173]	0.531 [-0.005, 1.119]	0.129 [0.067, 0.192]	0.546 [-0.079, 1.349]
Time investment	0	0.194 [0.128, 0.261]	0.016 [-0.691, 0.567]	0.213 [0.15, 0.276]	0.004 [-0.892, 0.696]

Note: Ninety-five percent confidence intervals in brackets based on 1,000 replications of the entire estimation procedure taking into account clustering at the municipality level.

of investments in which case we obtain a strong downward bias for the coefficient on material investments (0.11 and 0.129 under covariances A and B respectively instead of 0.594) and a strong upward bias for the coefficient on time investments as we would expect given the assumed structure of the covariance matrix of the errors. However, when we take into account the endogeneity of investments (*IV* columns) the bias almost completely disappears. For example, under the more conservative Covariance B we estimate the coefficient on material investments to be 0.546, when the true value is 0.594, and the coefficient on time investments to be 0.004, when the true value is zero. These results are strong evidence that our estimates do not suffer from weak instrument bias.³⁹

E. Complementarities between Inputs of the Production Function

We now go back to considering the implications of our estimates. The Cobb-Douglas specification implied by the data means that the inputs are complementary (Cunha et al. 2006, Heckman and Mosso 2014). However, in earlier versions of the paper, when we estimated the substitution elasticity, we always found it very close to 1. Given the metric we use for the latent factors, the return to investment is higher for children with better initial conditions. In panel A of Figure 2 we show

³⁹The full set of Monte Carlo results is presented in online Appendix Tables F.3 and F.4.

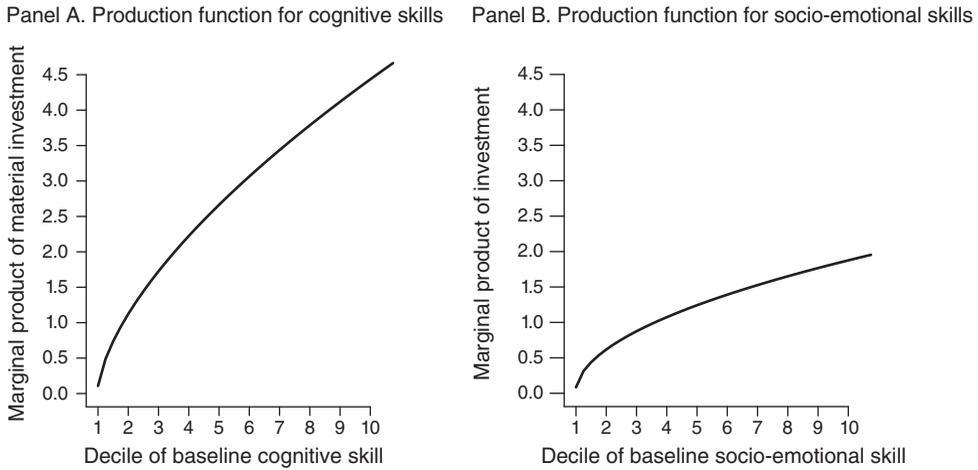


FIGURE 2. COMPLEMENTARITY BETWEEN INVESTMENTS AND BASELINE SKILLS

Notes: Panel A (panel B) is based on the estimates of the production function for cognitive skills (socio-emotional skills) reported in column 5 (1) of Tables 6 and 7, respectively. The figures 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 panel A and $\theta_{i,t}^S$ for panel B and holding all remaining inputs of the production function at their mean values across the sample.

the effect of a one standard deviation increase in material investment on cognitive skills and in panel B the effect of a one standard deviation increase in both material and time investments on socio-emotional skills. The y-axes are in standard deviation units of the outcome.

The complementarity of investments with initial conditions of the child may appear contradictory to the set of studies indicating that early interventions benefit low-achieving children the most (Bitler, Hoynes, and Domina 2014; Elango et al. 2016). However, one needs to allow for the differences in the populations concerned. Our intervention targets the 20 percent 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 programs such as Head Start in the United States. 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 socioeconomic distribution, diminishing returns could set in unless perhaps we design an intervention better attuned to higher ability children.

F. Implications for Longer-Term Outcomes

The results have implications for what to expect in the longer term. Under the assumption that the patterns of self-productivity and complementarities documented here remain the same at least in the medium term, the fact that the return to investments are complementary to the prior level of achievement implies that, if parents keep investing at the higher levels induced by the intervention, the skills of the treated children should continue to improve, subject of course to the impact of investments at later ages. And since ability and investments are complementary, future

investments would further increase the skills of the intervention group. However, the estimates also reveal fadeout between the baseline and the follow-up: the coefficient on past cognitive (socio-emotional) skills is about 0.7 (0.5) in the production function for cognitive (socio-emotional) skill for a time lapse of about 18 months. If such mean reversion continues beyond the ages that we consider and if parents revert to the level of activity in the control group, we can expect the impact of the intervention to become much smaller in the long run. In fact, as shown in Andrew et al. (2018), parents in our intervention had reduced their level of engagement to that of the control group when interviewed again two years after the end of the intervention. And consistent with the results here the impact of the intervention also faded.

Partial fadeout was also observed following the Jamaica intervention (Walker et al. 2005), although perhaps because the original effect was as large as 0.80 standard deviations, one-half of the original impact remained. This underscores two key lessons. First, we should not underestimate the challenges involved in scaling up successful small-scale efficacy trials, such as the Jamaica study, and in achieving comparable initial impacts. Second, we need to better understand what motivates parents to continue investing in children in the longer run and whether sustained intervention is needed to preserve and reinforce initial gains in such environments.

VI. Conclusion

Children from poor backgrounds accumulate developmental deficits from a very early age. 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 18-month-long early childhood intervention carried out in Colombia that promoted suitable parenting and stimulation to children between one and two years old at baseline. The intervention involved weekly home visits delivered by local women who had no prior specific knowledge of child development, but were trained as part of the intervention 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, language, and socio-emotional development, though impacts on the latter are smaller.

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

The estimates of the production functions also provide evidence of several important features of skill development among children below the age of 4. 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, but also implies mean reversion. Second, we find evidence of cross-productivity: the current stock of cognitive skills fosters the development of future socio-emotional skills, but the reverse does not seem to be the case. This result contrasts with that reported by Cunha, Heckman, and Schennach (2010), which finds socio-emotional skills to be important for the accumulation of future cognitive skills though at a different age. While the presence of siblings in the household reduces parental investments, it improves the cognitive and socio-emotional development of our subject child (who is the youngest child in the family in most cases), most likely through interactions and imitation, over and beyond the effect of parental investments.

Most importantly for the question addressed in this paper, as well as more broadly, our results show that investments help develop both cognitive and socio-emotional skills. The program increased investments substantially and our key conclusion is that it is this increase that led to the estimated impacts of the intervention on children's skills. Specifically, the intervention increased the cognitive development of the children by 0.115 log points and socio-emotional development by 0.087 log points. Our best estimates of the production functions imply that the increase in parental investments induced by the program account for all of the intervention impact on cognition and at least 66 percent of its impact on socio-emotional skills. There was no direct impact of the intervention. It is thus the parenting component, where the home visitor directly involves the mother (or main carer) in the stimulation activities and encourages ongoing engagement with the child based on the stimulating activities, that underlies its success.

Our study answers some important questions but raises many more, calling for further experimentation and analysis. We need to understand how to better target and treat the most disadvantaged of society. 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 long-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, an important lesson from our results is that involving parents in interventions is key to promoting child development in the short term. Going forward, it is crucial that we better understand how to ensure continued parental investments after the intervention has ended.

REFERENCES

- Agostinelli, Francesco, and Matthew Wiswall.** 2016. "Identification of Dynamic Latent Factor Models: The Implications of Re-Normalization in a Model of Child Development." NBER Working Paper 22441.
- Almlund, Mathilde, Angela Lee Duckworth, James Heckman, and Tim Kautz.** 2011. "Personality Psychology and Economics." In *Handbook of the Economics of Education*, Vol. 4, edited by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, 1–181. Amsterdam: Elsevier Science.
- Anderson, T. W., and Herman Rubin.** 1956. "Statistical Inference in Factor Analysis." In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 5, edited by Jerzy Neyman, 111–50. Berkeley, CA: University of California Press.
- Andrew, Alison, Orazio Attanasio, Emla Fitzsimons, Sally Grantham-McGregor, Costas Meghir, and Marta Rubio-Codina.** 2018. "Impacts 2 Years after a Scalable Early Childhood Development Intervention to Increase Psychosocial Stimulation in the Home: A Follow-Up of a Cluster Randomised Controlled Trial in Colombia." *PLoS Medicine* 15 (4): e1002556.

- Andrews, Donald W. K., and James H. Stock. 2005. "Inference with Weak Instruments." NBER Technical Working Paper 31.
- Attanasio Orazio P., Camila Fernández, Emla Fitzsimons, Sally M. Grantham-McGregor, Costas Meghir, and Marta Rubio-Codina. 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, Orazio, Emla Fitzsimons, Ana Gomez, Martha-Isabel Gutiérrez, Costas Meghir, and Alice Mesnard. 2010. "Children's Schooling and Work in the Presence of a Conditional Cash Transfer Program in Rural Colombia." *Economic Development and Cultural Change* 58 (2): 181–210.
- Attanasio, Orazio, Costas Meghir, and Emily Nix. 2019. "Investments in Children and the Development of Cognition and Health in India." NBER Working Paper 21740.
- Attanasio, Orazio, Costas Meghir, and Ana Santiago. 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, John E., Claire A. Bennett Freeland, and Mary L. Lounsbury. 1979. "Measurement of Infant Difficulty." *Child Development* 50 (3): 794–803.
- Bayley, Nancy. 2006. *Bayley Scales of Infant and Toddler Development*. 3rd ed. San Antonio, TX: Harcourt Assessment.
- Betancourt, Theresa S. 2015. "The Intergenerational Effects of War." *JAMA Psychiatry* 72 (3): 199–200.
- Bitler, Marianne P., Hilary W. Hoynes, and Thurston Domina. 2014. "Experimental Evidence on Distributional Effects of Head Start." NBER Working Paper 20434.
- Black, Maureen M., Susan P. Walker, Lia C. H. Fernald, Christopher T. Andersen, Ann M. DiGirolamo, Chunling Lu, Dana C. McCoy, et al. 2016. "Early Childhood Development Coming of Age: Science through the Life Course." *Lancet* 389 (10064): 77–90.
- Blundell, Richard, Alan Duncan, and Costas Meghir. 1998. "Estimating Labor Supply Responses Using Tax Reforms." *Econometrica* 66 (4): 827–761.
- Bound, John, David A. Jaeger, and Regina M. Baker. 1997. "Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak." *Journal of the American Statistical Association* 90 (430): 443–50.
- Campbell, Frances, Gabriella Conti, James J. Heckman, Seong Hyeok Moon, Rodrigo Pinto, Elizabeth Pungello, and Yi Pan. 2014. "Early Childhood Investments Substantially Boost Adult Health." *Science* 343 (6178): 1478–85.
- Campbell, Frances A., and Craig T. Ramey. 1994. "Effects of Early Intervention on Intellectual and Academic Achievement: A Follow-Up Study of Children from Low-Income Families." *Child Development* 65 (2): 684–98.
- Carneiro, Pedro, Karsten T. Hansen, and James J. Heckman. 2003. "Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College Choice." *International Economic Review* 44 (2): 361–422.
- Creech, Suzannah K., and Gabriela Misca. 2017. "Parenting with PTSD: A Review of Research on the Influence of PTSD on Parent-Child Functioning in Military and Veteran Families." *Frontiers in Psychology* 8: 1101.
- Cunha, Flavio, and James J. Heckman. 2008. "Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Journal of Human Resources* 43 (4): 738–82.
- Cunha, Flavio, James J. Heckman, Lance Lochner, and Dimitriy V. Masterov. 2006. "Interpreting the Evidence on Life Cycle Skill Formation." In *Handbook of the Economics of Education*, Vol. 1, edited by E. Hanushek and F. Welch, 697–812. Amsterdam: Elsevier Science.
- Cunha, Flavio, James J. Heckman, and Susanne M. Schennach. 2010. "Estimating the Technology of Cognitive and Non-Cognitive Skill Formation." *Econometrica* 78 (3): 883–931.
- Currie, Janet, and Douglas Almond. 2011. "Human Capital Development before Age Five." In *Handbook of Labor Economics*, Vol. 4B, edited by Orley Ashenfelter and David Card, 1315–1486. Amsterdam: Elsevier Science.
- Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall. 2014. "Household Choices and Child Development." *Review of Economic Studies* 81 (1): 137–85.
- Duflo, Esther, Rema Hanna, and Stephen P. Ryan. 2012. "Incentives Work: Getting Teachers to Come to School." *American Economic Review* 102 (4): 1241–78.
- Elango, Sneha, Jorge Luis García, James J. Heckman, and Andrés Hojman. 2016. "Early Childhood Education." In *Economics of Means-Tested Transfer Programs in the United States*, Vol. 2, edited by Robert Moffitt. Chicago: University of Chicago Press.
- Gelber, Alexander, and Adam Isen. 2010. "Children's Schooling and Parents' Behavior: Evidence from the Head Start Impact Study." *Journal of Public Economics* 101: 25–38.

- Gertler, Paul, James Heckman, Rodrigo Pinto, Arianna Zanolini, Christel Vermeersch, Susan Walker, Susan Chang, and Sally Grantham-McGregor.** 2014. "Labor Market Returns to an Early Childhood Stimulation Intervention in Jamaica." *Science* 344 (6187): 998–1001.
- Gorsuch, Richard L.** 1983. *Factor Analysis*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gorsuch, Richard L.** 2003. "Factor Analysis." In *Handbook of Psychology: Research Methods in Psychology*, Vol. 2, edited by John A. Schinka and Wayne F. Velicer, 143–64. Hoboken, NJ: Wiley.
- Grantham-McGregor, Sally, Yin Bun Cheung, Santiago Cueto, Paul Glewwe, Linda Richter, and Barbara Strupp.** 2007. "Developmental Potential in the First 5 Years for Children in Developing Countries." *Lancet* 369 (9555): 60–70.
- Grantham-McGregor, Sally, Christine A. Powell, Susan P. Walker, and John H. Himes.** 1991. "Nutritional Supplementation, Psychosocial Stimulation, and Mental Development of Stunted Children: The Jamaican Study." *Lancet* 338 (8758): 1–5.
- Heckman, James, Seong Hyeok Moon, Rodrigo Pinto, Peter Savelyev, and Adam Yavitz.** 2010. "Analyzing Social Experiments as Implemented: A Reexamination of the Evidence from the HighScope Perry Preschool Program." *Quantitative Economics* 1 (1): 1–46.
- Heckman, James J., and Stefano Mosso.** 2014. "The Economics of Human Development and Social Mobility." *Annual Reviews of Economics* 6: 689–733.
- Heckman, James, Rodrigo Pinto, and Peter Savelyev.** 2013. "Understanding the Mechanisms through Which an Influential Early Childhood Program Boosted Adult Outcomes." *American Economic Review* 103 (6): 2052–86.
- Helmers, Christian, and Manasa Patnam.** 2011. "The Formation and Evolution of Childhood Skill Acquisition: Evidence from India." *Journal of Development Economics* 95 (2): 252–66.
- Hu, Yingyao, and Susanne M. Schennach.** 2008. "Instrumental Variable Treatment of Nonclassical Measurement Error Models." *Econometrica* 76 (1): 195–216.
- Jackson-Maldonado, Donna, Virginia A. Marchman, and Lia C. H. Fernald.** 2012. "Short-Form Versions of the Spanish MacArthur-Bates Communicative Development Inventories." *Applied Psycholinguistics* 34 (4): 837–68.
- Kariger, P., E. A. Frongillo, P. Engle, P.M.R. Britto, S. M. Sywulka, P. Menon.** 2012. "Indicators of Family Care for Development for Use in Multicountry Surveys." *Journal of Health Popular Nutrition* 30 (4): 472–86.
- Knudsen, Eric I.** 2004. "Sensitive Periods in the Development of the Brain and Behavior." *Journal of Cognitive Neuroscience* 16 (8): 1412–25.
- Knudsen, Eric I., James J. Heckman, Judy L. Cameron, and Jack P. Shonkoff.** 2006. "Economic, Neurobiological, and Behavioral Perspectives on Building America's Future Workforce." *Proceedings of the National Academy of Sciences* 103 (27): 10155–62.
- Lu, Chunling, Maureen M. Black, and Linda M. Richter.** 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 (12): e916–22.
- Nielsen, Eric R.** 2015. "Achievement Gap Estimates and Deviations from Cardinal Comparability." Federal Reserve Board Divisions of Research & Statistics and Monetary Affairs Finance and Economics Discussion Series 2015-040.
- Putnam, Samuel P., Maria A. Gartstein, and Mary K. Rothbart.** 2006. "Measurement of Fine-Grained Aspects of Toddler Temperament: The Early Childhood Behavior Questionnaire." *Infant Behavior and Development* 29 (3): 386–401.
- Robin, Jean-Marc, and Richard J. Smith.** 2000. "Tests of Rank." *Econometric Theory* 16 (2): 151–75.
- Romano, Joseph P., and Michael Wolf.** 2005. "Stepwise Multiple Testing as Formalized Data Snooping." *Econometrica* 73 (4): 1237–82.
- Rothenberg, Thomas J.** 1971. "Identification in Parametric Models." *Econometrica* 39 (3): 577–91.
- Rubio-Codina, Marta, Orazio Attanasio, Costas Meghir, Natalia Varela, and Sally Grantham-McGregor.** 2015. "The Socioeconomic Gradient of Child Development: Cross-Sectional Evidence from Children 6–42 Months in Bogota." *Journal of Human Resources* 50 (2): 464–83.
- Schennach, Susanne M.** 2004. "Estimation of Nonlinear Models with Measurement Error." *Econometrica* 72 (1): 33–75.
- Schennach, Susanne M.** 2007. "Instrumental Variable Estimation of Nonlinear Errors-in-Variables Models." *Econometrica* 75 (1): 201–39.
- Staiger, Douglas, and James H. Stock.** 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65 (3): 557–86.
- Todd, Petra E., and Kenneth I. Wolpin.** 2006. "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility." *American Economic Review* 96 (5): 1384–1417.

- Walker, Susan P., Susan M. Chang, Christine A. Powell, and Sally M. Grantham-McGregor.** 2005. "Effects of Early Childhood Psychosocial Stimulation and Nutritional Supplementation on Cognition and Education in Growth-Stunted Jamaican Children: Prospective Cohort Study." *Lancet* 366 (9499): 1804–07.
- Walker, Susan P., Susan M. Chang, Christine A. Powell, Emily Simonoff, and Sally M. Grantham-McGregor.** 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, Susan P., Theodore D. Wachs, Sally M. Grantham-McGregor, Maureen M. Black, Charles A. Nelson, Sandra L. Huffman, Helen Barker-Henningham, et al.** 2011. "Inequality in Early Childhood: Risk and Protective Factors for Early Child Development." *Lancet* 378 (9799): 1325–38.