

*Annual Review of Economics***Early Childhood Development,
Human Capital, and Poverty****Orazio Attanasio,^{1,2,3,4,5} Sarah Cattan,^{2,6}
and Costas Meghir^{1,2,3,4,6,7,8}**

¹Department of Economics, Yale University, New Haven, Connecticut, USA;
email: orazio.attanasio@yale.edu, c.meghir@yale.edu

²Institute for Fiscal Studies, London, United Kingdom; email: sarah_c@ifs.org.uk

³Abdul Latif Jameel Poverty Action Lab (J-PAL), Cambridge, Massachusetts, USA

⁴National Bureau of Economic Research, Cambridge, Massachusetts, USA

⁵FAIR—Centre for Experimental Research on Fairness, Inequality and Rationality, NHH
Norwegian School of Economics, Bergen, Norway

⁶Institute of Labor Economics (IZA), Bonn, Germany

⁷Centre for Economic Policy Research, London, United Kingdom

⁸Institute for Evaluation of Labour Market and Education Policy (IFAU), Uppsala, Sweden

Annu. Rev. Econ. 2022. 14:853–92

First published as a Review in Advance on
June 1, 2022

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-092821-053234>

Copyright © 2022 by Annual Reviews.
All rights reserved

JEL codes: J24, O15, I24, I25, I3

Keywords

early childhood development, human capital, inequality, poverty,
education, economic development

Abstract

Children’s experiences during early childhood are critical for their cognitive and socioemotional development, two key dimensions of human capital. However, children from low-income backgrounds often grow up lacking stimulation and basic investments, which leads to developmental deficits that are difficult, if not impossible, to reverse later in life without intervention. The existence of these deficits is a key driver of inequality and contributes to the intergenerational transmission of poverty. In this article, we discuss the framework used in economics to model parental investments and early childhood development and use it as an organizing tool to review some of the empirical evidence on early childhood research. We then present results from various important early childhood interventions, with an emphasis on developing countries. Bringing these elements together, we draw conclusions on what we have learned and provide some directions for future research.

**ANNUAL
REVIEWS CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

1. INTRODUCTION

The links between poverty and human capital development are multifaceted and complex. Gaps in many dimensions of development between poorer and richer children emerge early and persist through the life cycle. In turn, the importance of human development in determining adult outcomes means that the link between family income and human capital plays a substantial role in explaining the intergenerational transmission of poverty observed to varying degrees across the world (Björklund & Salvanes 2011, Black & Devereux 2011, Alesina et al. 2021). At the macro level, the link between poverty and human capital development is central to understanding dynamics of productivity, growth, and inequality.

The developmental deficits of children living in poverty have been documented in several contexts. **Figure 1** shows the relationship between a family wealth index and a well-established measure of vocabulary (the Peabody Picture Vocabulary Test) in nationally representative samples of 5-year-olds in three developing countries. Although the exact shape of this relationship varies across contexts, the figures shown are a disheartening reminder of the developmental disadvantage that children living in poverty experience from the youngest age.

These large inequalities in early development are so concerning because they are known to map onto later ones. **Figure 2** shows that children's vocabulary at age 5 is highly predictive of their vocabulary at age 15 (panel *a*), which in turn is highly predictive of their educational attainment at age 22 (panel *b*). Although the data in these two panels are drawn from different samples, a simple imputation exercise demonstrates how strongly gaps in education measured in adulthood can be traced back to early childhood skills. While these figures obviously do not necessarily reflect a causal effect of early childhood development on later outcomes, they do reflect the long-lasting influence that early childhood inequalities have over the life cycle.

Three series of the *Lancet* journal have addressed deficits in child development associated with poverty.¹ Although these articles focused on developing countries, the issues they raise are




Figure 1

Child language development at age 5 by family wealth quintile in (a) Ethiopia, (b) India, and (c) Peru. The figure plots the average percentile of children in the distribution of language development at age 5 for each quintile (Q) of the family wealth distribution within each country. Language is measured by the Peabody Picture Vocabulary Test (PPVT), based on raw PPVT scores standardized for age (in months) using local linear regressions. Wealth is measured by an index constructed by latent factor analysis of indicators for whether the child's parents own the house; whether the household has access to electricity, drinkable water, and kerosene or gas fuel for cooking; whether the household possesses some durable goods (radio, fridge, bike, television, motor car or tractor, pump, phone, and sewing machine); and whether the house's walls, roof, and floors are made of raw natural material (e.g., wood or soil but not concrete). N = 1,860 for Ethiopia, 1,851 for India, and 1,903 for Peru. Data from the Young Lives Study (Boyden 2021).

¹Readers are referred to Grantham-McGregor et al. (2007), *Lancet* (2007), and all papers in the *Lancet* series.

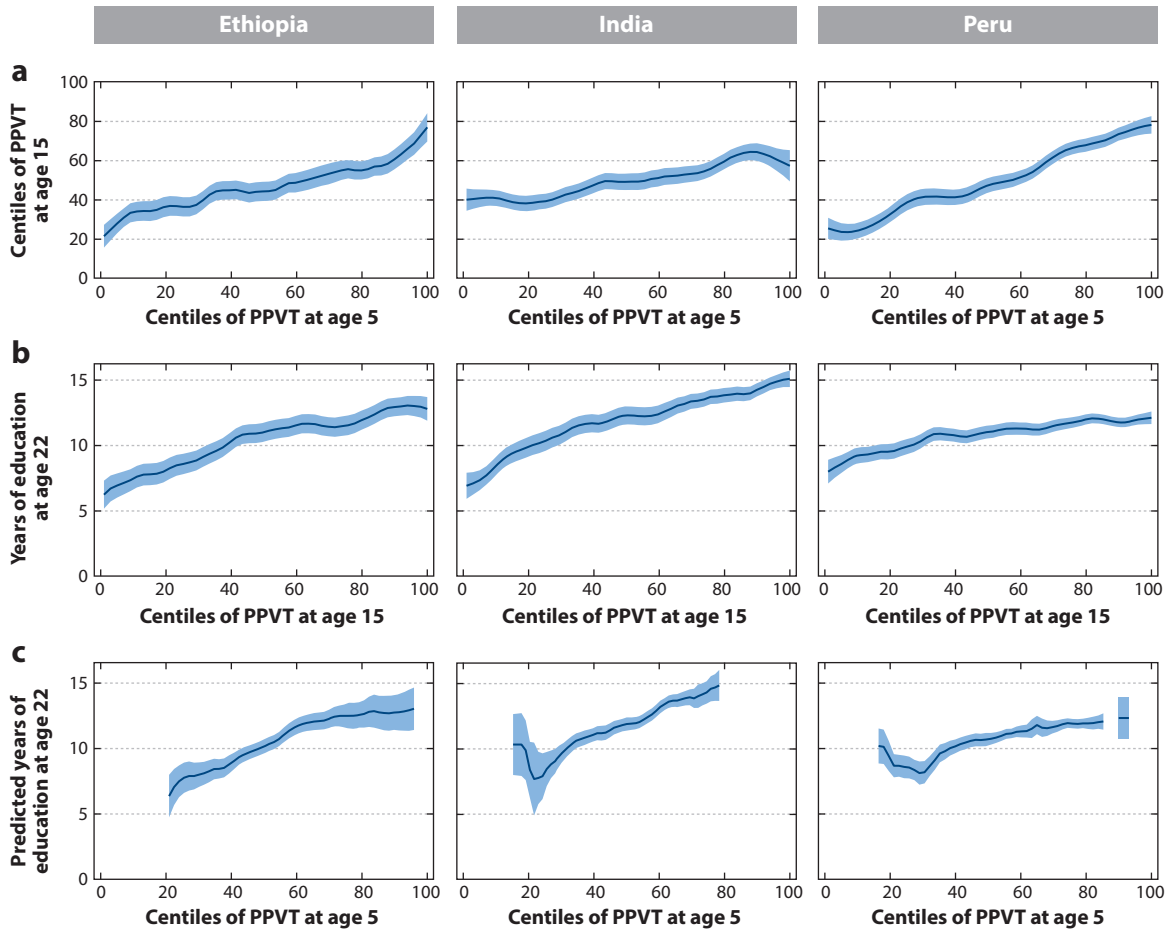


Figure 2

The predictive power of early childhood skills for later skills and educational attainment. Language is measured by the Peabody Picture Vocabulary Test (PPVT). (a) Rank-rank relationship between PPVT at age 5 and PPVT at age 15. (b) Relationship between PPVT rank at age 15 and years of education at age 22. (c) Predicted relationship between PPVT rank at age 5 and years of education at age 22, using relationships depicted in panels a and b. Samples sizes for the younger and older cohorts are, respectively, 809 and 456 for Ethiopia, 980 and 866 for India, and 946 and 639 for Peru. Data from the Young Lives Study (Boyden 2021).

universal. However, understanding how we should address these developmental deficits is less well understood. In order to do that, we need a deep understanding of the fundamental reasons for these deficits.

The aim of this article is to discuss the link between poverty and human development. We pay particular attention to the mechanisms through which this link operates and to what these imply for the design of appropriate interventions aimed at reducing poverty through investments in human capital. Given the importance of the early childhood period for adult human capital, we focus on this period of the life cycle, though much of our theoretical discussion about the human development process is relevant to other periods of life. We review the interdisciplinary nature of the knowledge base, and we emphasize the insights that can be gained from interpreting the evidence through the lens of an economic framework, touching upon parental behavior in relation to the way children are brought up.

We therefore start the article by presenting a simple economic model of household behavior and child development. The model presupposes that households maximize lifetime utility, which depends on private consumption as well as the human capital of their offspring (or a function of that).² This household optimization problem is subject to technological constraints (the production function for child human capital), informational constraints, and a budget constraint. The production function models the dynamic process through which endowments and inputs (chosen by members of the household) determine the evolution of child human capital. The information constraints relate to the beliefs that parents have about how human capital is produced and the effectiveness/productivity of the time and resources (investments) that they devote to the upbringing of their children.

The model is helpful for clarifying the possible channels through which poverty can affect human development. Within the same economic environment, the socioeconomic gradient in human development can be due to gaps in parental investments and initial endowments; and inequalities in parental investments can be explained by inequalities in endowments, beliefs, financial resources, and preferences. We use this framework to structure our review of the evidence. First, considering the process of human capital development, we discuss advances in the specification, identification, and estimation of production functions for child development. What do we know about the inputs that matter? What are the challenges to identifying the role of inputs? How should the dynamics and interactions among inputs be modeled, and why does it matter for our understanding of the link between poverty and human development?

We then turn to the drivers of parental investments. We review the evidence looking at the socioeconomic gradient in parental investments and discuss the various explanations proposed in the literature. Do poor households invest less in their children because they have fewer financial resources, different preferences for these investments, different perceptions about their returns, and/or different beliefs over what constitute appropriate child-reading practices?

This evidence has some important implications for policy aimed at breaking the link between poverty and human capital development, though there are still important unknowns about how we should address the developmental deficits of poor children. A standard view, in particular among economists, has long been that those deficits are the results of a combination of preferences and resources. The focus on resources, however, ignores the fact that investments in early childhood do not necessarily demand important commitments of resources. In contrast, the child development literature emphasizes the importance of child stimulation through language interaction and simple games that can be based on common household materials. Simply making conversation with a child, reading them a book, and involving them in household and play activities can have large effects on cognitive development. From this perspective, deficits in child development are mostly due to a lack of information and knowledge about the child-rearing process, though this literature also recognizes that the stress induced by poverty could prevent parents from engaging in these activities. Financial constraints may become more important later in childhood, but if developmental deficits are already present in the early years, even then the returns to later investments may be low.

This lack of knowledge and information, often reinforced by cultural conventions about raising children, may be a key reason for deficits in child development, and it underlies the logic guiding parenting interventions in early childhood. These include, for example, the ABCderian program (Campbell & Ramey 1994), the Nurse Family Partnership (Olds et al. 2004, 2014,

²It is straightforward to consider households that care about child human capital more generally, and we need not constrain ourselves to Becker's altruistic model (Becker & Tomes 1986).

2019), and the Jamaica Home Visit (JHV) program (Grantham-McGregor et al. 1991). Parenting interventions take the form of guiding parents to stimulate their children, often starting right after birth. Importantly, many of these interventions do not advocate for large increases in costs or even in time spent with children, but rather for a change in parents' interactions with them: for example, eliminating negative actions such as corporal punishment, introducing positive actions such as expressions of praise and affection, and introducing stimulating play activities based around the daily chores in a household.

In what follows, we describe such interventions and present results from a number of experiments, including some designed and analyzed by the authors. The aim of this discussion is to illustrate the possibilities and the difficulties that need to be addressed if this approach becomes standard in preventing developmental deficits. The large and sustained human development impacts of parenting interventions evidenced in the literature pose an important challenge to traditional economic thinking about the nature of interventions needed to remedy developmental gaps among low-income children. Indeed, this literature suggests that, if it becomes possible to change beliefs, and in some cases social norms, so that parent-child attachment improves and low-income families engage in stimulating activities, then the need for intervention will abate in the next generation. Based on empirical results on the complementarities of child development across ages (Cunha et al. 2010, Attanasio et al. 2020e), it is also highly likely that an early successful intervention will increase the returns to later ones, including interventions to improve the quality of preschool and later levels of education. Although the results are drawn from developing countries, the key issues they emphasize do not differ fundamentally from those that need to be addressed in pockets of poverty in wealthier countries, although the exact content and implementation model of the intervention would likely require adaptation in practice.

2. HUMAN CAPITAL DEVELOPMENT AND POVERTY: A GENERAL FRAMEWORK

This section presents a general framework to help clarify the link between poverty and the process of human capital development within and across generations. The model characterizes the dynamic process through which human capital develops throughout childhood, starting in the very first years of life, and embeds this process within a model of altruistic household behavior. This approach helps outline the mechanisms through which poverty persistently affects human capital development across generations, and it suggests some levers that might be available to policy makers seeking to improve child outcomes.

The framework considers the developmental trajectory of an individual child with some genetic endowment at conception until adulthood.³ The child is exposed to a variety of factors, including the physical environment where they live; their interactions with parents, other family members, possibly teachers, and other children; and a certain number of inputs, such as nutrition, health care, and formal schooling. These factors, along with shocks, determine the development of the child since conception.

The process of human development can be summarized by way of a set of production functions, in which a vector of inputs at a given age determines an outcome or a particular level of human development at a later age (for an early application, see, e.g., Leibowitz 1974; for a discussion of theoretical and empirical issues, see Todd & Wolpin 2003). The seminal papers of Cunha & Heckman (2008) and Cunha et al. (2010) have addressed many of the challenges associated with

³As we mention below, the model can be extended to consider several children, incorporating both fertility choices (and the quality-quantity trade-off) and the problem of resource allocation across different children.

identifying and estimating this model and have given the dynamic production function center stage in the analysis of child development with observational and experimental data. In its most general form, the production function can be written as

$$\mathbf{H}_{i,a} = F_a(\mathbf{H}_{i,a-1}, \mathbf{H}_{i,a-2}, \dots, \mathbf{H}_{i,0}, \mathbf{X}_{i,a}, \mathbf{X}_{i,a-1}, \mathbf{X}_{i,a-2}, \dots, \mathbf{X}_{i,0}, \boldsymbol{\epsilon}_{i,a}), \quad 1.$$

where $\mathbf{H}_{i,a}$ is a vector that represents the human development of child i at age a , and $\mathbf{X}_{i,a}$ is a vector of observable inputs. $H_{i,0}$ represents the child's initial condition or endowment at birth, which will be the result of genetics and in utero nutrition. The function F_a represents the process of human development. It is indexed by the child's age to reflect the fact that the productivity of inputs can change with age. $\boldsymbol{\epsilon}_{i,a}$ is a vector of unobserved inputs, including shocks such as incidents of ill-health. These may be observed by the parents or teachers but not by the researcher.

The model above is agnostic about the nature of inputs relevant to child development. It is also very general in allowing the whole history of inputs and prior levels of human development to enter the production function. When taking this model to the data, researchers will need to make decisions about which inputs and how many lags to include in the production function. Often, these decisions will be influenced by the data available. Indeed, many of these inputs may be difficult to observe or may not be observed at all, which creates issues with identification. As we discuss below, these decisions can have important implications for the estimates of the model.

The lack of complete data on the relevant inputs to the process in Equation 1, which are reflected in $\boldsymbol{\epsilon}_{i,a}$, makes it difficult to estimate the effects of inputs that are observed. The difficulty arises from the fact that some inputs of interest are chosen by multiple agents with a specific stake in the outcome under study. These agents may include parents, teachers, siblings, and the children themselves. Depending on the child's age, the relative importance of the different agents will change. For example, in the early years, we would expect the primary carers of the child (and potentially the child's preschool teachers) to have most agency. As the child grows up, the decisions of the child and of their peers would become more influential for the child's human capital development. As we primarily focus on the early childhood period in this article, we concentrate our discussion on the drivers of parental behavior.

During this period of the life cycle, the family plays a prominent role, and there is broad consensus that inequalities in children's outcomes originating in early childhood are in large part determined by inequalities in how stimulating, nurturing, and safe the children's home environments are. There is ample correlational and causal evidence showing that parental behavior—in terms of both the warmth and closeness of parent–child relationship and the type of activities that parents engage in with their children—are key for the cognitive and socioemotional development of children during this age. There is also much evidence that poor parental mental health, parental stress, and parental conflict are important risk factors for the development of young children (especially for their emotional and behavioral developments). This could have a direct effect on child development or an indirect one, through lower parental engagement and attachment with the child. Whereas all these inputs are strongly correlated with family income, whether financial resources have a role per se is disputed. In fact, many of the activities promoted by parenting interventions require very little financial resources.

Regardless of who makes investment decisions in the household, these choices are likely to be related to elements of the developmental process that are not directly observed by the researcher. In other words, these variables are endogenous. A possible and useful strategy for the identification of such causal links is to model the endogenous variables—and therefore, in this context, to model the behavior of the agents making the relevant decisions. Such an approach requires the definition of the drivers of parents' and possibly teachers' behavior and choices. This might help identify

variables that can affect these choices without affecting the child's developmental outcomes directly. As is often the case, the identification of such variables can be problematic. However, without a strategy of this kind, it may be difficult, if not impossible, to identify the marginal effect of the endogenous inputs on child development.

Below we sketch a stylized model of parental investment behavior to outline key sets of drivers. Following a long tradition of models of altruistic parental behavior (e.g., Becker & Tomes 1986), we assume that parents in household i maximize a function that depends on child development \mathbf{H}_α^i at some final age α and on the path of their own consumption, C_t^i , $t = 1, \dots, \alpha - 1$.⁴ While parents choose C_t^i directly, \mathbf{H}_α^i is the outcome of a production function whose arguments are the level of child development, \mathbf{H}_t^i , $t = 1, \dots, \alpha - 1$; the path of parental investment, \mathbf{X}_t^i ; and, possibly, some other factors \mathbf{Z}_t^i and ϵ_t^i , with the former being observable and the latter unobservable. The problem can then be summarized as

$$\max_{(C_t^i, \mathbf{X}_t^i)_{t=1}^\alpha} \sum_{t=1}^{\alpha} \beta^t U_i(C_t^i) + V_i(\mathbf{H}_\alpha^i) \quad 2.$$

$$\begin{aligned} \text{subject to} \quad \mathbf{H}_t^i &= \tilde{\mathbf{f}}_i(\mathbf{H}_{t-1}^i, \mathbf{X}_t^i, \mathbf{Z}_t^i, \epsilon_t^i | \Omega_i), \\ A_{t+1}^i &= (1 + r_t)A_t^i + Y_t^i - C_t^i - \mathbf{p}_t \mathbf{X}_t^i, \end{aligned}$$

where Y_t^i is income, \mathbf{p}_t is the price of parental investments, and A_t^i are assets that can be used to move resources over time at the interest rate r_t . We stress that $\tilde{\mathbf{f}}_i(\cdot)$ is the production function as perceived by parents, that is, conditional on their information set Ω_i . This information may not be fully accurate, and so the perceived production function may not correspond to the true production function outlined in Equation 1 if parents have distorted beliefs over its inputs or parameters.

This model of parental behavior is highly stylized and omits a number of features, some of which we discuss in subsequent sections.⁵ However, even in this simple version, it is clear that parental investment decisions depend on parents' preferences and three types of constraints. The first one is technological: As discussed above, the process of development follows a particular technology that maps inputs into outputs (human development). The second one is financial: Investment may be costly, and parents may be limited in their ability to borrow to finance their investment. The third type of constraint is informational: Parents make decisions subject to their perceptions of their child's level of human development and of the production function, both of which may correspond to reality only imperfectly.

While the budget constraint introduces a mechanical link between family poverty and child development, poverty may affect parental investment behavior through all other components of the model. Children born in poorer households may start with lower endowments or initial conditions \mathbf{H}_0^i . Depending on the dynamics of the production function and on parental preferences, this

⁴It is possible that parents derive utility over the full trajectory of their child's human capital development; or they could derive utility over their children's adult outcomes (e.g., earnings or well-being), which could be modeled as a function of their child's human capital at age α or as a function of their child's trajectory of human capital development.

⁵Among others, the framework assumes a unitary model of decision making, whereas each parent may have a distinct set of preferences over investments and consumption. Moreover, the model assumes the family has a single child, whereas there could be multiple children with different initial conditions. We return to this in Section 4.2.

lower endowment may have a long-lasting effect on the child's outcomes either directly (through the production function) and/or by affecting parental investments. There is also evidence suggesting that parents from poorer backgrounds may have different information sets compared to parents from richer backgrounds, they may have different preferences over their children's outcomes, and they may even have different production functions. All these systematic differences across the income distribution may be responsible for creating a stubborn link between poverty and human development. We discuss them in turn in the rest of this article, starting with key features of the production function in the next section.

3. THE PRODUCTION FUNCTION FOR CHILD DEVELOPMENT

The production function for child development characterizes the process of human capital accumulation over time. To be useful, the general production function in Equation 1 has to be made specific, but doing so raises several important challenges. First, it is necessary to establish the dimensions of the vector \mathbf{H} . Moreover, one needs to specify the set of relevant inputs, their potential persistence, and how they interact among themselves within and across periods. We discuss these topics in turn.

3.1. Dimensions of Human Capital

For a long time, economists have considered human capital a low-dimensional variable that would enter the production function used in many models. In the simplest models, production is a function of one-dimensional human capital, and individuals are heterogeneous in their human capital endowment. More sophisticated models considered two types of human capital, usually skilled and unskilled, which play different roles in the production process—for example, in the way they interact with other factors of production like capital and raw materials. In most of these models, these different types of human capital are mutually exclusive: Individuals are endowed with one type or another (e.g., Katz & Murphy 1992).

Over the past 20 years, the conceptualization of human capital in both micro- and macroeconomic models has radically changed. Human capital is now increasingly recognized as a multidimensional object, and this allows for a much richer characterization of heterogeneity across individuals.⁶ While levels of human capital may be correlated across dimensions, this correlation is far from perfect. As exemplified in the important work of Heckman and coauthors on the General Educational Development (GED) testing program, individuals with the same average level of cognitive skill but different levels of noncognitive skills may end up with vastly different outcomes (Heckman & Rubinstein 2001, Heckman et al. 2014).

Modeling human capital as a multidimensional object is important because different dimensions of human capital play an important, and different, role in determining later outcomes. A large literature finds evidence of robust associations between various dimensions of human capital and a range of adult outcomes, such as educational attainment, labor market outcomes, criminal engagement, healthy behaviors, teenage pregnancy, and marital stability (e.g., Heckman et al. 2006, Almlund et al. 2011, Lundberg 2017, Heckman & Karapakula 2019, Carneiro et al. 2022, Cattan et al. 2022). An interesting study by Berniell & De la Mata (2016) reports estimates of returns to cognitive and socioemotional skill in 10 Latin American countries, stressing that these skills might have different relevance depending on the occupational structure of the labor market.

⁶Readers are referred to the surveys by Acemoglu & Autor (2011) and Deming (2017).

Several papers also look at the returns to different types of socioemotional skills on a variety of outcomes. For example, Papageorge et al. (2019) find that a higher level of externalizing behavior⁷ leads to lower educational attainment but a higher wage in the labor market (conditional on education). Returns to these skills in the labor market may also change over time, for example, in response to changes to the production function process (Borghans et al. 2014, Deming 2017).

Importantly for our discussion, different dimensions of human capital and different types of skills can interact with each other in the development of human capital over the life cycle. These interactions, labeled cross-productivity by Cunha et al. (2006), have been documented empirically in several papers. Cunha & Heckman (2008) and Cunha et al. (2010) find evidence that higher levels of early noncognitive skills boost the development of later cognitive skills. They interpret these findings as suggesting that a child who is better able to focus might be better able to exploit educational opportunities and accumulate cognitive skills. Evidence of cross-productivity has also been found across other domains of human capital. For example, Attanasio et al. (2020c) find evidence that higher levels of cognitive skill at ages 1–2 lead to higher levels of socioemotional skills at ages 3–4. Attanasio et al. (2020e) find that better health at age 5 leads to higher cognitive skills at age 8.

In **Table 1**, we compare the main features of a selected set of papers estimating production functions for child development. Whereas most papers focus on the joint development of cognitive and socioemotional skills and/or look at their simultaneous impact on adult outcomes, very recent applied work in economics also distinguishes between socioemotional dimensions. Using data from the British cohort studies, several papers distinguish between internalizing and externalizing behaviors. For example, Attanasio et al. (2020d) study the evolution of these two behaviors along with cognitive skills, showing that they interact with each other in dynamically complex ways and have long-run effects on labor market outcomes. Moroni et al. (2019) model the development of externalizing and internalizing behaviors in middle childhood, allowing the productivity of a set of parental inputs to vary both across the distribution of children's socioemotional skills in early childhood and across the input distribution.

From an economic point of view, the multiple dimensions of human capital are important if their relative price in the labor market changes over time, which implies that they cannot be aggregated into one Hicks-aggregate human capital. These relative price changes may occur because of technological change and/or in response to changes in the supply of different skills. Several papers find evidence that the return to noncognitive skills may have increased over the past decades in many economies. To explain this trend, Deming (2017) focuses on the role of individuals' ability to work in teams as a skill that is complementary to cognitive skills. He attributes the increase in the return to this skill to the increase in specialization in production (whereby individual workers contribute to a limited number of tasks). Because the outputs of these very specialized tasks have to be combined to produce final outputs, the return to noncognitive skills, such as the ability to work in teams, might increase and become as high as the return to cognitive skills, if not more so. Hence, if the relative valuation of alternative skills changes in the labor market (e.g., because of technical change), the discussion on the dimensionality of human capital becomes central to understanding the changing occupational and wage structures and, consequently, inequality.

⁷Externalizing behaviors are often contrasted with internalizing behaviors, with the former including behaviors related to aggression and hyperactivity and the latter including behaviors related to anxiety and withdrawal.

Table 1 Comparison of main features of production functions for child development in a selected set of papers

Paper	Country & data set	Outcomes	Inputs	Dynamics	Functional form	Treatment of endogeneity
Agostinelli & Wiswall (2020)	United States, National Longitudinal Survey of Youth	Cognition, schooling, and earnings	Parental investment	One lag of the outcome	Translog	Instrumental variables
Agostinelli et al. (2020)	United States, Early Childhood Longitudinal Study–Kindergarten Class of 1998–1999	Cognitive and noncognitive skills	Home and classroom investments	One lag of the outcome	Translog	Instrumental variables
Attanasio et al. (2020a)	Colombia, Acio TU RCT	Cognition, health, and socioemotional skills	Parental investment and skills	Two lags of the cognition, health, and socioemotional skills that vary with child's age	CES, translog, and Cobb–Douglas	Control function approach
Attanasio et al. (2020c)	Colombia, 2014 micronutrient supplementation and psychosocial stimulation RCT	Cognition and socioemotional skills	Parental material and time investments	One lag of cognitive and socioemotional skills	Cobb–Douglas	Instrumental variables
Attanasio et al. (2020d)	England, Scotland, and Wales; 1970 British Cohort Study	Internalizing and externalizing socioemotional skills, absolute mobility indexes, and cognition	Parental socioemotional skills	Socioemotional and cognitive skills at age 5 for the production function at age 10, and at ages 5 and 10 for the production function at age 16	Cobb–Douglas	Controlling for observed characteristics
Attanasio et al. (2020e)	India, Young Lives Survey	Cognition and health	Parental material investment, cognition, and health	One lag of cognition and health	CES	Control function approach
Attanasio et al. (2017)	Ethiopia and Peru, Young Lives Survey	Cognition and health	Parental material investment, cognition, and health	One lag of cognition and health	CES	Control function approach
Caucutt et al. (2020)	United States, Panel Study of Income Dynamics–Child Development Supplement	Cognition	Parental material and time investments	One lagged outcome	Cobb–Douglas	Instrumental variables

(Continued)

Table 1 (Continued)

Paper	Country & data set	Outcomes	Inputs	Dynamics	Functional form	Treatment of endogeneity
Chaparro et al. (2020)	United States, Infant Health and Development Program	Cognition and maternal care quality	For cognition: quantities and qualities of child care provided by the mother, the program, and others For maternal care quality: quantities of care provided by the mother	One lagged outcome	Cobb-Douglas	Instrumental variables for the maternal care quality and controlling for observed characteristics for cognition
Cunha & Heckman (2008)	United States, Children of the National Longitudinal Survey of Youth 1979	Cognitive and noncognitive skills and earnings	Parental investment	One lag of cognitive and noncognitive skills	Cobb-Douglas	Instrumental variables
Cunha et al. (2010)	United States, Children of the National Longitudinal Survey of Youth 1979	Cognitive and noncognitive skills and earnings	Parental investment and cognitive and noncognitive skills	One lag of cognitive and noncognitive skills	CES	Control function approach
DelBoca et al. (2014)	United States, Panel Study of Income Dynamics–Child Development Supplement	Cognition	Parental active and passive time investment	One lag of the outcome	Cobb-Douglas	Controlling for observed characteristics
Del Bono et al. (2020)	United Kingdom, Millennium Cohort Study	Cognitive and noncognitive skills	Child's and mother's cognitive and noncognitive skills	One lag of the outcome alone and interacted with lagged child's and parental skills	Cobb-Douglas	Controlling for observed characteristics
Helmers & Patnam (2011)	India, Young Lives Survey	Cognition, health, and noncognitive skills	Parental material investment	One lag of the outcome	Cobb-Douglas	Instrumental variables

Abbreviations: CES, constant elasticity of substitution; RCT, randomized controlled trial.

3.2. The Form of the Production Function and Its Implications

It is now widely recognized that environmental factors in the very early years of life have long-lasting impacts on individual development and adult outcomes.⁸ There is also evidence that inputs interact with one other to produce future levels of human capital. Together with the dynamics of the production function, the nature of the interactions between inputs in the same period determines the extent of dynamic complementarities between investments made in different periods (Cunha et al. 2006).

Even though these are crucial issues, an appropriate and exhaustive characterization of the dynamic properties of human development and of substitution patterns between inputs is still unavailable. We discuss how the current literature has handled these issues and highlight the challenges to be tackled by future research.

3.2.1. Input substitution within and across periods. As shown in **Table 1**, researchers have used various functional forms for Equation 1 that allow for more or less flexible interactions between inputs. Several studies specify a Cobb–Douglas (CD) specification, which implies a unit elasticity of substitution among different inputs (e.g., Cunha & Heckman 2008, DelBoca et al. 2014, Attanasio et al. 2020c). The CD specification imposes very strong restrictions on the substitutability among different inputs, an assumption that may be particularly severe when considering several inputs. Some studies, such as those by Cunha et al. (2010) and Attanasio et al. (2020e), generalize the CD specification to consider a constant elasticity of substitution (CES) specification, which encompasses the CD form as a special case. Whereas the CES is less restrictive than the CD specification, it still imposes strong assumptions on the pattern of substitutability of different inputs. In particular, it implies that any pair of inputs has the same elasticity of substitution.

In turn, some papers have sought to relax this assumption by considering a nested CES specification, which allows different groups of inputs to have different elasticities of substitution. Caucutt et al. (2020) use a nested CES in which they define parental investment as a function of parental time inputs and market child care. The aggregate investment then interacts with other inputs on the basis of a CD production function to generate child outcomes. Another example is provided by Attanasio et al. (2017), who estimate a nested CES function to model health and cognitive development at ages 8, 12, and 15 using data from Peru and Ethiopia. They strongly reject the restrictions implied by the standard CES for a sample of Ethiopian children collected as part of the Young Lives Survey.

One disadvantage of both the CES and the nested CES is that the relationship between the relevant variables and the outcomes of interest can be highly nonlinear, therefore involving considerable econometric challenges. Moreover, the requirement that the function be concave restricts the substitution patterns. An attractive alternative is to consider a translog production function in which the output is modeled as a second-order polynomial in the (log of) prior achievement, investment, and other background variables. Such a specification preserves linearity in parameters while allowing a considerable amount of flexibility. This functional form is used by Attanasio et al. (2020a).

These different patterns of substitutability across inputs can have important implications for the process of child development and, therefore, for the design of policies. Attanasio et al. (2017) present some simulations plotting the impulse response function of human development to an exogenous change in parental investment at age 5. In the first scenario, this shift in investment

⁸Readers are referred, for instance, to the discussions by Cunha et al. (2006), Almond & Currie (2011), and Elango et al. (2016).

occurs while holding baseline health constant. In the second scenario, this shift is accompanied by an increase in the child's initial health level at age 5. This exercise aims to capture how the dynamics of the interactions between different inputs (in this case, parental investment and initial health) vary depending on the exact specification of the production function (here, CES or nested CES).

This exercise shows that, as expected, an early shift in investments has a large positive effect on the evolution of cognition over time. However, both the magnitude of the impact and the interaction between the shift in investment and initial health status are, in some contexts, quite different across the two specifications of the production function. This exercise also shows that the extent to which increasing initial health boosts the effect of the early shift in investments on cognition over time also varies between the CES and nested CES production functions. These differences come from the fact that (a) the estimates of the marginal product of investment at different ages are different when one allows for the flexibility implied by the nested model, and (b) the complementarities among health, cognition, and investments also differ between the CES and nested CES specifications.

3.2.2. Dynamics. Although the prevailing consensus is that the process of human capital formation is inherently dynamic, the exact nature of these dynamics is still relatively unknown. The dynamics are important in two ways. On the one hand, the parameters that characterize the production function can change substantially with age. On the other hand, the persistence of lagged skills can be very different across the dimensions of skill considered. Equation 1 is most general in that it allows human capital at age a to be a function of the whole history of inputs and previous levels of human capital. However, identifying such function empirically involves important data requirements (in terms of both the frequency with which data are collected and the length of the panel), which exceed the features of most available data sets.

Because we let the production function in Equation 1 be age specific, the process is intrinsically nonstationary. This means that the productivity of certain inputs may vary with age. Moreover, the exact nature of this nonstationarity may determine the existence of windows of opportunity for intervention aimed to bolster the development of children with early developmental delays. This is the case if certain dimensions of human development in one period have direct persistent effects on later outcomes (i.e., over and beyond their effects on skills in the next period working through self-productivity).

As shown in **Table 1**, it is common practice in empirical studies to assume that the human capital accumulation process follows a first-order Markov process—that is, outcomes at age a depend only on outcomes at age $a - 1$, given that we condition on current inputs. Formally, if $\mathbf{H}_{i,a}$ is the vector that represents the level of development of child i in its various dimensions at age a , then a first-order Markov process describing the evolution of $\mathbf{H}_{i,a}$ with age can be expressed as

$$\mathbf{H}_{i,a+1} = \mathbf{g}_a(\mathbf{H}_{i,a}, \mathbf{X}_{i,a}, \boldsymbol{\epsilon}_{i,a+1}), \quad 3.$$

where $\mathbf{X}_{i,a}$ is a vector of observable (exogenous and endogenous) variables determined at age a . More generally, we note that in Equation 3 we let the function $\mathbf{g}_a(\cdot)$ vary with age. The important point we want to stress here is that, conditional on $\mathbf{H}_{i,a}$ (and the other inputs $\mathbf{X}_{i,a}, \boldsymbol{\epsilon}_{i,a+1}$), the previous level of development, reflected in $\mathbf{H}_{i,a-j}, j > 0$, is assumed to be irrelevant for $\mathbf{H}_{i,a+1}$.⁹

⁹One could think of the unobservable term $\boldsymbol{\epsilon}_{i,a+1}$ as being composed of a time-invariant component, $\boldsymbol{\eta}_i$, and of a time-varying component, $\mathbf{v}_{i,a+1}$, which is possibly correlated over time. Even if $\mathbf{v}_{i,a+1}$ is not serially correlated, the presence of $\boldsymbol{\eta}_i$ creates a source of persistence in the unobservable in the production function, to be distinguished from state dependence working through $\mathbf{H}_{i,a}$.

Such a model, while convenient from the analyst's point of view, might be unable to explain certain empirical observations, such as the fade-out and subsequent reemergence of the impacts of certain interventions. There are several ways in which this Markov assumption can be relaxed. First, it is possible that the dynamics of the process are more complex than those described in Equation 3. A simple extension, for instance, would be to consider a model of the type

$$\mathbf{H}_{i,a+1} = \tilde{\mathbf{g}}_a(\mathbf{H}_{i,a}, \mathbf{H}_{i,a-1}, \mathbf{X}_{i,a}, \boldsymbol{\epsilon}_{i,a+1}). \quad 4.$$

An alternative possible violation of the Markovian assumption embedded in Equation 3 is that there could exist some critical age, possibly very early in the life cycle, when development is key for subsequent development in a way that is not summarized by subsequent attainment levels. Define such age as α^* . It is then possible that the right model is

$$\mathbf{H}_{i,a+1} = \hat{\mathbf{g}}_a(\mathbf{H}_{\alpha^*}, \mathbf{H}_{i,a}, \mathbf{X}_{i,a}, \boldsymbol{\epsilon}_{i,a+1}), \quad \forall a > \alpha^* - 1. \quad 5.$$

One important fact, which has been observed in several studies and which speaks against the simple first-order Markov assumption, is the long-run impact of some early child development interventions on adult outcomes, even when no impacts are apparent at intermediate ages (Bailey et al. 2017, 2020). One well-known example is the Perry Pre-School Program, in which early impacts of the intervention on intelligence quotient (IQ) scores disappeared only to reemerge in other domains, including socioemotional skills (Heckman et al. 2010). Another example is the JHV program, which provided psychosocial stimulation to children aged 9–24 months. In this case, large impacts on cognition measured at the end of the intervention decreased in magnitude over time and were no longer statistically significant when children were 7–8 years of age (Grantham-McGregor et al. 1997). However, the intervention was found to have improved a variety of adult outcomes, including earnings and criminal behavior, measured about 20 years after the end of the intervention (Gertler et al. 2014).

Ultimately, the distinction among Equations 3–5 is an empirical matter. However, distinguishing among them is hindered by the scarcity of appropriate longitudinal data covering a sufficiently long period. Moreover, the measures used to capture development throughout childhood are often different across age periods, making longitudinal links difficult and dependent on the specific anchoring chosen (Cunha et al. 2010, Agostinelli & Wiswall 2016). In practice, the dynamics of the empirical specifications of equations such as Equation 3 or 4 are driven by the frequency at which data are observed, as reflected in **Table 1**.

A few recent papers have modeled the process of child development using data from developing countries and focusing on the dynamics of the process. In particular, Attanasio et al. (2020a) use a high-frequency data set, which contains information on child development and other outcomes, collected on a sample of Colombian children observed from birth to 7 years, roughly at an annual frequency.¹⁰ They model child development in three dimensions: cognition, socioemotional skills, and health. They find evidence of several interactions across skills, and in the case of cognition they find that incorporating more than one lag of development is important for explaining future outcomes.

More generally, Attanasio et al. (2020a) find that the production function changes considerably over time, both in the impact of different inputs and in the level of persistence each dimension exhibits. In the case of cognition, for example, the level of persistence increases considerably with age, and the productivity of parental investment is significant until age 4 and then declines considerably. In the case of socioemotional development, the productivity of parental investment

¹⁰The data are an unbalanced panel of five waves, containing children of different ages.

becomes important after age 4. For health, persistence is very high from very early on in the life cycle.¹¹

Although there is still too little evidence to determine whether the patterns found by Attanasio et al. (2020a) are specific to the context they study or are more general, the dynamic properties of the process of human capital development have crucial implications for determining when and for how long it is best to intervene to promote human development. For example, if persistence in cognition is lower in the first two years of life than it is in the next two, but the marginal productivity of parental investment is higher in earlier periods than it is later, this suggests that an optimal policy would be to intervene very early but sustain investments until an age at which depreciation (or fade-out) is unlikely. To date, still too little is known about the dynamics of the process of human capital development across short subsequent periods of the life cycle. We see this as an important priority for research moving forward.

4. THE DETERMINANTS OF PARENTAL INVESTMENTS

The model we sketched in Section 2 provides a framework for understanding the factors that influence child development, and ultimately it can offer a structural approach for identifying the causal effects of investments in children on adult outcomes. The production function represented in Equation 1 is central to the model, and its characterization is key for establishing what policies could improve child development among disadvantaged children. The main difficulty with such a characterization is to establish and quantify the causal pathway from certain inputs chosen by parents or other agents to the development of the child. As shown in the last column of **Table 1**, the literature has adopted various approaches to this issue, ranging from ignoring it to making different assumptions about the endogeneity of parental choices. Most approaches, however, rely on a model of parental behavior, which in turn motivates the choices of a set of instrumental variables (or the construction of a control function).

In this section, therefore, we discuss the main determinants of parental investment. Within the relatively stylized structure we sketched in Section 2, we identified the main drivers of parental behavior and investment as (a) resource constraints, (b) tastes and preferences, and (c) the (perceived) process of child development. We now discuss them in turn.

4.1. Resource Constraints

As set out in the model in Section 2, a first reason parents make different investments in their children's human capital is that they have different financial resources available to do so. The lower the resources available (or the higher the price of investments), the lower the probability that parents will invest in their children's human capital. If the imperfection arises from the inability to transfer resources from the future to the present (i.e., liquidity constraints, which are not fully explicit in the model we presented), a positive effect could be obtained by increasing current parental income and/or subsidizing investments (e.g., Becker & Tomes 1986, Dahl & Lochner 2012).¹²

Evidence on the causal effect of increasing family income, in the sense of a pure income effect, is rare to find. This is because most reforms that have been used by researchers to identify the effect of family income on children's outcomes (or investments in their human capital) not only increase family income but also change other inputs that may also affect child development. In

¹¹We note that health-specific investment data are not available in this data set.

¹²Carneiro et al. (2021) examine the importance of timing of parental income over the life cycle for child outcomes.

particular, most cash transfer policies, such as PROGRESA in Mexico and other similar programs (many of which have been rigorously evaluated in many low-income contexts), make transfers conditional on parental investments in children, such as schooling or health care (Fernald et al. 2008, Attanasio et al. 2010). Welfare reforms, such as the Earned Income Tax Credit (EITC), which have been used by Dahl & Lochner (2012) to study the impact of family income on children's outcomes, increase family income as well as maternal labor supply. The policy therefore changes the allocation of parental time and the quality of the care provided to the child, which may affect children's outcomes holding family income constant (Agostinelli & Sorrenti 2018).

Although the vast literature on cash transfers has mostly considered conditional cash transfers,¹³ important insights can be obtained from the handful of papers that have evaluated the impact of unconditional cash transfers. Macours et al. (2012), for instance, show that an unconditional cash transfer in Nicaragua led to improvements in early childhood development: Children in households that were randomized into an intervention called Atención a Crisis had significantly higher levels of development than children in the control group 9 months after transfers were started. Furthermore, these impacts persisted 2 years after the program had been discontinued and the transfers ended. While this evidence is consistent with an important role for financial resources in determining children's outcomes, the authors present evidence suggesting that other program features, such as the social marketing that accompanied the transfers, or the fact that transfers were made to women, were likely to be important in explaining the results.¹⁴

Another issue that has received recent attention is the fact that the lack of appropriate financial resources, coupled with a stressful environment and a lack of social support system, may impose a considerable amount of strain on parents and prevent them from performing even simple parenting tasks to stimulate their child's development.¹⁵ Within this context, a set of programs that consist in transferring assets (and possibly training) to ultrapoor households are particularly interesting. These programs, first tried in Bangladesh and subsequently replicated and validated in a number of Low- and Middle-Income Countries (LMICs), have received considerable attention. The studies reviewed by Banerjee et al. (2015) have shown that such programs have had, in most places, a considerable impact on individual incomes and more generally well-being. In doing so, these programs might help households to escape poverty traps. Unfortunately, information on child development has not been collected in studies evaluating these programs. Such an analysis would be particularly pertinent, as many of the asset and skill transfers that have been analyzed are targeted to women and, therefore, are more likely to change women's labor supply and bargaining power within the household.

4.2. Parental Tastes and Preferences

The simple model above posits that parental preferences depend on child development, H , and own consumption, C . In the most basic form of the model, heterogeneity in preferences over H and C across households could lead to observing different investment behaviors across households. Any correlation between such heterogeneity in preferences and family income could create a

¹³Readers are referred to the surveys by De Walque et al. (2017), Bastagli et al. (2020), and Molina-Millan et al. (2020) on the long-term effects of such transfers. The literature on conditional cash transfers is extensive; Fiszbein et al. (2009) provide a review.

¹⁴Maluccio & Flores (2005) study the same program in Nicaragua, and Paxson & Schady (2010) ask a similar question for a program in Ecuador.

¹⁵On the effect of stress on different types of performance, readers may consult Mani et al. (2013).

socioeconomic gradient in investment (which would contribute to the socioeconomic gradient in child development).

To be made more realistic, however, the model could be enriched in a variety of ways. First, one could distinguish between different types of investment (e.g., material and time investments) and could allow parents to have preferences over performing certain activities with their children. Second, parents may also be heterogeneous in their preferences for different domains of child development. For example, some parents may care about their children's cognitive development more than they care about other dimensions of development, whereas others may value socioemotional development and health more. This type of heterogeneity could help explain the differential patterns of investment observed across different demographic groups, defined by socioeconomic status and cultural or religious norms.¹⁶ Some authors argue that this heterogeneity in preferences could be driven by heterogeneity in beliefs over the returns to various skills. For example, Kohn (1963) argues that mothers from lower socioeconomic backgrounds have stronger preferences for the development of their children's socioemotional skills because they believe obedience and conformity have high labor market returns.

The recent work by Doepke & Zilibotti (2017, 2019) and Doepke et al. (2019) looks at how parents choose their parenting style based on their specific set of preferences and on their perception of the returns that different parenting styles will have for child development (and future outcomes). In their model, two types of preferences are important in driving such behavior: altruism, which determines how much parents care about their children's utility, and paternalism, which determines how much parents care about their children's actions in ways that potentially conflict with the children's own preferences. More generally, heterogeneity in preferences (and/or beliefs about the developmental process, which we discuss below) might be behind the remarkable differences in parental investment across households from different backgrounds. Dotti Sani & Treas (2016), for instance, report that across many countries more educated parents (whose time in the job market should be more valuable) spend more time with their children than less educated ones.

The model we have discussed so far assumes a single child and omits fertility choices, which in developing countries are an important consideration, in particular when thinking about the quantity-quality trade-offs that poor parents might face. The model could be extended in this direction without much difficulty. Such considerations, however, would introduce a number of other important dimensions to the parents' problem. Indeed, for families with several children, preferences will also have to incorporate a taste for equality among children that parents might have, another dimension that can be added to the basic model. How resources (and eventual outcomes) are distributed across children will play an important role in driving investment decisions. If children are born with (or develop in the early years) different skills and endowments, parents might face a trade-off between efficiency (i.e., maximizing the total level of development of their children) and possible equity concerns. As discussed by Almond & Mazumder (2013), parental investment strategies might attempt to compensate for perceived differences in initial conditions or strategically reinforce them, depending on the features of the perceived process of child development and on their preferences. Marginal returns to investments and, consequently, chosen levels of investment may also be affected by the characteristics of the child. Differences in ability across children may therefore affect how investments are allocated across and within households (e.g., Behrman et al. 1982, Aizer & Cunha 2012).

¹⁶For example, Lynd & Lynd (1929, 1937) reported that working-class mothers ranked strict obedience as their most important child-rearing goal more frequently than mothers from higher socioeconomic backgrounds did.

How parents distribute resources and investment across different children matters. Giannola (2021) reports that, in many countries, inequality among siblings accounts for a substantial fraction of total inequality in earnings among individuals. Furthermore, he shows that, whereas average outcomes across siblings in various dimensions decline with family size, the best outcome among siblings does not vary with family size. This result is robust to a number of considerations and indicates that parental investment has an important role in determining inequality among siblings.¹⁷

A number of recent papers empirically explore these questions in developing countries. For example, Adhvaryu & Nyshadham (2016) use data from Tanzania and exploit variation in initial conditions induced by a randomly allocated pre-birth intervention. They show that parents make larger health investments in children who are born with stronger endowments as a result of the intervention. This is consistent with the idea that parents reinforce birth endowments through their investments, in accordance with the findings reported by Giannola (2021). Using data on twins from China, Yi et al. (2015) argue instead that the family acts as a net equalizer in response to early health shocks across children. Berry et al. (2020) explore this question with a lab-in-the-field experiment aimed at identifying parents' preferences for the total amount of development across siblings as well as their trade-offs with inequalities in outcome or inputs. They show that although parents do care about average earnings, they also have a strong preference for equality in inputs. They do not find evidence that parents care about equality in outcomes.

Finally, gender-specific preferences may also be important in some contexts: Parents may allocate resources and time based on preferences for the gender of the child, even superseding the returns to such allocations.¹⁸ In the simple model considered so far, the decision units are the parents, considered as a monolithic block with well-defined preferences. However, maternal and paternal preferences might be different, which implies that parental investment decisions will depend on both sets of preferences as well as on the nature of the decision process within the household. There is vast empirical evidence that is consistent with this hypothesis. For example, Thomas (1994), Hoddinott & Haddad (1995), Doss (2006), and Schady & Rosero (2008) show that income controlled by women is associated with higher expenditures on food as well as with higher-quality nutrition (e.g., higher in protein) and, often, improved child outcomes. Macours & Vakis (2010) show nonexperimental evidence on the positive impact of mother's seasonal migration on children's cognitive development. Lundberg et al. (1997) and Ward-Batts (2008) present quasi-experimental evidence from the United Kingdom to argue that income is more likely to be spent on clothing for women and children when it is controlled by women than when it is controlled by men.

Models of intra-household decision making have received considerable attention in recent years, with their implications for parental investment decisions being the focus of more recent research. The theoretical underpinnings of the collective model with public goods are provided by Blundell et al. (2005), who focus on expenditure for children as the public good.¹⁹ In a recent paper, Almås et al. (2021) model parental investment in a semistructural fashion, using data elicited to measure bargaining power within couples in Tanzania.²⁰ In their application, parental behavior

¹⁷Giannola (2021) also finds that parents in the slums of a city in India reinforce differences in early endowments.

¹⁸Fertility choices have also been shown to depend on the gender of existing children (Butcher & Case 1994).

¹⁹The collective model of Chiappori (1988) has set the foundation for much of this work. For reviews, readers are referred to Bourguignon & Chiappori (1992) and Chiappori & Meghir (2015).

²⁰In this paper, the authors use the measure of bargaining power within a couple designed by Almås et al. (2018).

depends, as in the model sketched above, on tastes, bargaining power, and beliefs about the process of child development.

An important and underresearched question is how parental preferences are formed and, eventually, how new social norms about investing in children evolve in society. An interesting paper by Wang et al. (2022) hypothesizes that parents take a social reference point as a norm. This typically would be based on their observations; in the case of the paper in question, it is outcome based, namely, the height of other children. If a shock then shifts the reference point, it could change parental perceptions of this outcome and shift their behavior to achieve it, increasing investment and leading to a new norm. Their empirical work is based on data from a nutritional experiment in Guatemala, which substantially increased the height of treated children. This shifted the reference point of parents of children born later, leading to increased nutritional investments and improved height. A broader understanding of how social norms and parental perceptions can be changed is likely to be of central importance in designing policies to improve child development.

The discussion so far should make it clear that parental preferences can be complex, and they can reflect different sources of heterogeneity that lead to different types of parenting practices and, ultimately, differences in child development. Attempts to measure directly the drivers of individual behavior, including preferences, can be valuable for establishing the causal links between parental investment and child development. Some of the papers cited above undertake this strategy and are important in paving the way for further work in this direction.

4.3. Informational Constraints

In standard economic models, parents are assumed to be rational and to know the production function for human development. If that is the case, variation in parental investment can only be explained by variation in tastes and variation in resources, which drive variation in inputs (observed or unobserved).²¹ The past decade has seen a burgeoning of evidence pointing to the importance of relaxing this assumption to consider the role of informational imperfections in determining investments in human capital.

The literature discusses two particular sources of information friction. The first one is around parents' perceptions of their children's abilities. The second one is around parents' perceptions of the technology of human capital formation. Both may be incorrect, and the extent to which they are distorted may be correlated with socioeconomic status. This, in turn, could be another mechanism through which poverty is linked to lower human capital investments (and hence lower children's outcomes).

4.3.1. Parental perceptions of children's abilities. Parents might misperceive the ability of their children, either in absolute or in relative terms. This issue has been studied in both developed countries (see, for instance, Kinsler & Pavan 2021) and developing ones. For the latter, Dizon-Ross (2019) shows in the context of Malawi that misperceptions are more common among the poorest parents, and that providing information to parents can change parental choices substantially. Parental misperceptions may also be compounded by misperceptions about the returns to different educational choices (by parents or youths). Several studies have used observational data to show that students' beliefs about their own abilities predict their decisions, such as choice

²¹It may well be that there are unobserved components in the production function (say, total factor productivity) that differ across households. We can conceive of these as unobserved inputs.

of college major or college dropout.²² These findings also complement a recent information experiment performed in Mexico by Bobba & Frisnacho (2020), who tested predictions about the differential roles of the mean and variance of beliefs on educational decisions.

4.3.2. Parental perceptions about the production function. Another potential source of information imperfections, which is particularly relevant in our context, relates to parental beliefs about the effects that parental investment have on child development. The salience of these imperfections can be different for parents from different socioeconomic backgrounds (see, e.g., Lareau 2003, Putnam 2015). For example, while all parents might care equally about the development and well-being of their children, low-income parents might not be aware (or as aware as high-income parents) of the importance that some specific activities, such as talking to and interacting in specific ways with a young child, might have for their development. As Lareau (2003) argues, this may be because they believe that child development follows a natural growth process, thus underestimating the extent to which the brain is malleable and shaped by early stimulation.

To explain the relationship between parental investment and family income, Caucutt et al. (2017) propose a model in which parents misperceive the child development production function, and especially the usefulness of early years investment. Other types of studies elicit direct information on parental beliefs about the process of child development. Cunha et al. (2013, 2020) design innovative instruments that allow direct elicitation of quantitative measures of individual perceptions and find evidence of such misperceptions in a sample of disadvantaged mothers in the United States.²³

Attanasio et al. (2019) further develop these methods to measure parental beliefs about the productivity of investments in the context of an early parenting intervention for low-income families in Colombia (we return to this in Section 5). They show that mothers in their sample underestimate the productivity of parental investment substantially. Moreover, they find that mothers tend to view parental investment as being more useful for children with relatively low levels of development than for children with higher levels of development. Estimates of the production function in the same sample suggest that parental investment complements baseline levels of skills, which contradicts the latter set of maternal beliefs. Finally, Attanasio et al. (2019) show that their measures are meaningful: Parental investment is positively correlated with parental beliefs about its productivity.²⁴

Despite mounting evidence about the importance of parental investments for child development, there are still important gaps in our understanding of the drivers of parental behavior and of the link between family income and the quantity and quality of parental investments. Though the role of financial constraints cannot be underestimated, the evidence reviewed so far suggests a strong role for the lack of information and knowledge about the process of child development in explaining deficits in child development.

This lack of knowledge and information, often reinforced by cultural conventions about child-rearing, underlie the logic of early childhood parenting interventions. These focus on demonstrating good practice for stimulating children and strengthening the way parents interact with

²²Readers are referred to Chevalier et al. (2009), Arcidiacono et al. (2012), Stinebrickner & Stinebrickner (2012), and Stinebrickner & Stinebrickner (2014).

²³Distorted beliefs can also be important for educational choices beyond the early years, as discussed by Boneva & Rauh (2018) and Attanasio et al. (2020b), among others.

²⁴In a recent paper, Giannola (2021) combines data from a survey on parental investments and parental beliefs with data from a lab-in-the-field experiment in India showing that parents do not seem to have a strong taste for equality in outcomes among siblings.

their children. In the next section, we discuss the large literature evaluating these interventions and interpret the evidence through the lens of the economic framework underpinning this article. The discussion illustrates the possibilities and difficulties that need to be addressed if this approach becomes standard in preventing developmental deficits. Although the results are drawn from developing countries, the key issues do not differ fundamentally from those that need to be addressed in pockets of poverty in wealthier countries, although some of the features may differ in practice.

5. EVIDENCE ON PARENTING INTERVENTIONS

There is broad consensus that inequalities in children's outcomes originating in early childhood are in large part determined by inequalities in family environments and home stimulation. It has also been argued that economic circumstances might be only one, and not even the most important, determinant of child development; that the family plays a prominent role; and that stimulation is possible and effective even with limited resources. Against this background, parenting interventions are a promising direction for mitigating or even reversing such early developmental inequalities between poorer and richer children. These policies support caregivers to enrich the home environment and the quality of their interactions with children in order to provide greater stimulation and strengthen the emotional bond between children and their parents.

Many early childhood parenting interventions have been tried in developing countries over the last few decades. One of the best known is Care for Development (CfD), which has been promoted extensively by the WHO and UNICEF.²⁵ An intervention somewhat similar in spirit to CfD is the JHV program, which was first implemented in the 1970s and, unlike others, was rigorously and continuously evaluated; it culminated in a seminal study described in a series of papers (e.g., Grantham-McGregor et al. 1991, 1997). The JHV was not the first home-visiting intervention ever implemented in developing countries,²⁶ but it is one of the few interventions that have repeatedly followed up with the participants to evaluate long-term impacts. Quite remarkably, the latest study on this intervention, by Gertler et al. (2014), shows that the labor market earnings of the treated group improved by 25% by the time children were 22 years old.

As we discuss in this section, these results are remarkable in that they defy the notion that remediating the developmental deficits of poor children must require substantial financial resources. Indeed, some of these interventions (as the ones we focus on in this section) have been shown to be low cost compared to interventions implemented in the United States in the early years that have been shown to deliver long-term impacts (such as the Perry Pre-School Program, the ABCdarian program, and Family Nurse Partnership). Moreover, and importantly for their scalability, they can be implemented by local, nonspecialist staff. This last aspect is relevant not only for the financial cost of running such interventions, which is obviously key in low-resource settings, but also for the interventions' ability to change behavioral patterns in disadvantaged communities.

²⁵There are several reviews of stimulation interventions; for reviews of such interventions in LMICs, readers may consult, for instance, Baker-Henningham & López Bóo (2010), Engle et al. (2011), and Richter et al. (2019).

²⁶In addition to the CfD program, two early childhood interventions for children suffering from malnutrition implemented in Colombia in the 1970s also inspired the JHV program. One was an experimental study conducted in Cali between 1971 and 1974, which evaluated the impact of high-quality preschool program providing integrated health, nutritional, and educational activities (McKay et al. 1978). The second one was a nutritional and psychosocial stimulation program implemented between 1973 and 1976 in Bogotá, Colombia (Super et al. 1990). This study randomized children into four groups: one that received the nutritional component only, one that received the psychosocial stimulation program only, one that received both, and a fourth one that received neither.

Despite the promise that these interventions hold to promote child development in the early years among disadvantaged families, there remain important questions about how they can be successfully adapted and targeted to yield significant and long-term impacts, especially at scale. At what age should this type of program start? How long and how frequently should parents be solicited by the program? What dimensions of development (language, cognition, socioemotional skills) should they mostly target at different ages? How should these interventions be delivered at scale so that they do not lose their effectiveness? And how can we ensure that the short-term impacts of these interventions do not fade out over time? All these are important questions that researchers and policy makers alike should consider in order to refine the design of these promising interventions and implement them at scale.

In what follows, we review the body of evidence surrounding the JHV program and the interventions that were modeled after it and experimentally evaluated in Bangladesh, China, Colombia, and India. We focus specifically on the JHV program because it has been adapted in its content and mode of delivery in a variety of contexts. Its curriculum, now registered under the name of The Reach up Early Childhood Parenting Program, has been used as the core of several efficacy and effectiveness trials in, among other places, Bangladesh, China, Colombia, and India, and as the basis of a universal early childhood program in Peru. Although what follows is not an exhaustive review of the literature on parenting interventions, we argue that the evidence surrounding this single-parenting program (which shares many common features with other such programs) provides a remarkable opportunity to reflect on the factors behind its effectiveness in the short and the long term. In turn, this analysis can help inform how to target and deliver the programs better so that they achieve maximum cost-effectiveness and scalability in the future.

5.1. The Reach Up Early Childhood Parenting Program

The Reach Up Early Childhood Parenting Program (henceforth Reach Up) is based on the JHV intervention, described for instance by Grantham-McGregor et al. (1991). The program works by trying to build a positive relationship between parents and children and by strengthening parenting skills through a number of home visits occurring at regular intervals for an extended period, between 9 months and 2 years in the available studies. Each visit starts with a review of the activities introduced in the previous weeks. The home visitor then introduces a new set of activities for the parent to perform with the child during the coming week and discusses how these can be included in daily routines. Each activity is designed to address a separate developmental domain, such as cognitive (puzzles), language (stories, songs, books, etc.), and motor skills. The activities are supported by materials, including picture books, story cards, and toys, typically made with waste materials, such as plastic bottles and cloth, so as to help affordability and encourage the caregiver to produce their own.

The intervention is a highly structured and somewhat prescriptive curriculum that can be delivered by a well-trained home visitor who does not necessarily have formal qualifications. The activities to be performed during the visits are precisely described in the curriculum manual the visitor uses. Such activities become progressively more complex as children grow. Although each visit is mapped to a specific age, measured in weeks, during the training the visitors are encouraged to use earlier or later visits to match the level of complexity of the activities to the developmental stage of the child. Moreover, the program requires very low levels of resources. No materials are given to the parents, although some are left in the household for a week and then exchanged with the new materials used for progressing with the intervention. In other words, there is no element of subsidy implicitly related to the program. All of this means that the program is feasible in low-income settings and potentially scalable.

Since its initial use in Jamaica in the 1970s and 1980s, the program has been adapted for and trialed in various cultural contexts, including Colombia, India, Bangladesh, and China. Focusing on those interventions that have been experimentally evaluated, **Table 2** compares several of their features as well as impacts on child development. Although all these interventions are based on the same curriculum and hence on the same developmental and pedagogical approach, there are important differences in terms of the population they target and the ways in which they deliver the Reach Up curriculum. For example, the Jamaican and Bangladeshi interventions focus on undernourished or severely undernourished (stunted) children (Grantham-McGregor et al. 1991, Hamadani et al. 2006, Nahar et al. 2012, Tofail et al. 2013, Hamadani et al. 2019). In contrast, the Colombian, Indian, and Chinese interventions focus on low-income populations but do not require children to show signs of undernourishment. Specifically, in Colombia, the intervention was offered to families who were eligible to receive the Colombian conditional cash transfer program known as Familias en Acción (Attanasio et al. 2014). In India, one intervention targeted migrants living in the slums in Cuttack (Andrew et al. 2019), and another targeted poor children living in rural areas of Odisha (Grantham-McGregor et al. 2020).²⁷ In China, the interventions targeted children in an underdeveloped rural area in northwestern China (Sylvia et al. 2020, 2022).

Another key difference between these different adaptations of the same curriculum is in the way the curriculum was delivered. The JHV and several of its adaptations (Colombia, Cuttack, one of the Chinese trial, and two of the four Bangladeshi trials) implemented it via home visits. The other two Bangladeshi trials implemented it via group sessions. The Odisha study is the only one to implement, within the same trial, both home visits and group sessions so that their relative effectiveness could be compared. One of the Chinese trials (Sylvia et al. 2022) delivered the curriculum in a parenting center, where caregivers were invited for one-on-one parenting training sessions and also encouraged to bring their children to engage in free play, socialize with other children, and participate in organized activities. We return to this below when discussing issues of scalability.

Although the program was never delivered by child development specialists, the background and qualifications of the home visitors differed depending on location and context: health paraprofessionals in Jamaica, female local community leaders in Colombia,²⁸ community workers associated with an existing and well-established NGO in India (Pratham), local women and health workers in Bangladesh, and parenting trainers from the National Health Commission in China. The choice of who delivered the intervention was directly related to the program's scalability and sustainability: the emphasis on local women with no specific qualifications, but a strong training on the curriculum delivery, clearly solves any problem of scarcity of human resources and has the potential to make the program culturally more acceptable within their communities.

5.2. Heterogeneity in Short-Term Impacts Across and Within Studies

The experimental evaluations of these interventions point to their overall effectiveness in improving children's cognitive development (see **Table 2**, column 10). Nevertheless, impact sizes [measured in units of standard deviations (SD) of the control group] do vary, sometimes quite dramatically, across studies of the same core intervention. For example, in the Colombia study,

²⁷The authors of Grantham-McGregor et al.'s (2020) publication are Sally Grantham-McGregor, Akanksha Adya, Orazio Attanasio, Britta Augsborg, Jere Behrman, Bet Caeyers, Monimalika Day, Pamela Jervis, Reema Kochar, Perna Makkar, Costas Meghir, Angus Phimister, Marta Rubio-Codina, and Karishma Vats.

²⁸These women, known as Madres Lideres, are in charge of the local administration the Colombian conditional cash transfer program called Familias en Acción.

Table 2 Summary of experimentally evaluated interventions following the Jamaica Home Visit or Reach Up program

Country	Intervention	Delivery mode of psychostimulation	Duration	Target population	Age at baseline	RCT	Sample size	Horizon	Impacts (only reported if $p < 0.1$)	Reference
Jamaica (1986 to 1987)	Psychosocial stimulation and/or nutritional supplementation	Weekly home visits delivered by community health aides	24 months	Stunted children (WAZ < -2)	9 to 24 months old	Individual-level (1: stim; 2: nutrition; 3: stim + nutrition; 4: control)	129 children	Endline	COG: 0.91 stim, 0.76 nutrition LM: 0.49 stim, 0.59 nutrition LANG: 0.59 stim H&E: 0.80 stim PRF: 0.63 stim, 0.53 nutrition	Grantham-McGregor et al. (1991)
Colombia (2010 to 2012)	Psychosocial stimulation and/or nutritional supplementation	Weekly home visits by local female community leaders (Madre Líderes)	18 months	Beneficiaries of Colombia's conditional cash transfer (Familias en Acción)	12 to 24 months old	Cluster-level (1: stim; 2: nutrition; 3: stim + nutrition; 4: control)	96 towns 1,420 children	Endline 2 years after intervention end	COG: 0.26 RL: 0.22 No significant impact	Atanasio et al. (2014) Andrew et al. (2018)
India (Cutback (2013 to 2015))	Psychosocial stimulation	Weekly home visits by community workers	18 months	Migrants living in urban slums	12 to 24 months old	Cluster-level	54 slums 378 children	Endline	COG: 0.35 RL: 0.22 EL: 0.19	Andrew et al. (2019)
India (Odisha) (2015 to 2017)	Psychosocial stimulation or nutritional supplementation	Weekly home visits or group sessions	24 months	Children living in rural areas	7 to 16 months old	Cluster-level (1: stim via home visits; 2: stim via groups; 3: nutrition; 4: control)	192 villages 1,400 children	Midline and endline	COG: 0.32 home visits, 0.28 group sessions LANG: 0.24 home visits, 0.30 groups sessions	Grantham-McGregor et al. (2020)
	Enhanced preschool experience	Training of the preschool (Anganwadi) staff	18 to 24 months	Children living in rural areas	31 to 40 months old	Cluster-level (children from trial above rerandomized into enhanced or basic preschool)	192 villages 1,400 children	Trial in progress	NA	NA
Bangladesh (Monohordan)	Psychosocial stimulation	Weekly group sessions and twice-weekly home visits by local women	12 months	Children with moderate and severe under-nutrition (WAZ < -2)	6 to 24 months old	Cluster-level	20 community nutrition centers 193 children	One year after intervention end	COG: 0.27 SE: 0.28 COOP: 0.41 VOC: 0.30	Hamadani et al. (2006)

(Continued)

Table 2 (Continued)

Country	Intervention	Delivery mode of psychostimulation	Duration	Target population	Age at baseline	RCT	Sample size	Horizon	Impacts (only reported if $p < 0.1$)	Reference
Bangladesh (Dhaka)	Psychosocial stimulation and/or nutritional supplementation	Twice-weekly parenting sessions (1–3 dyads) delivered by local women	6 months	Children with severe under-nutrition (WAZ < -3)	6 to 24 months old	Individual level (1: stim; 2: nutrition; 3: stim + nutrition; 4: clinic control, 5: hospital control)	504 children	Endline	COG: 0.33 stim + nutrition WAZ: 0.27 stim + nutrition LAZ: 0.30 stim + nutrition	Nahar et al. (2012)
Bangladesh (Monohardam)	Psychosocial stimulation	Weekly group sessions delivered by health workers	9 months	Children with iron deficiency anemia and children who were neither anemic nor iron deficient	6 to 24 months old	Cluster level	30 villages 434 children	Endline	COG: 0.38	Tofail et al. (2013)
Bangladesh (Narsingdi) (2015 to 2016)	Psychosocial stimulation	Weekly group sessions delivered by health workers	12 months	Children with moderate and severe under-nutrition (WAZ < -2)	5 to 24 months old	Cluster level	90 community clinics 687 children	Endline	COG: 1.13 LANG: 0.87 FM: 1.02 SE: 0.89	Hamadani et al. (2019)
China (2014 to 2015)	Psychosocial stimulation	Weekly home visits by local trainers	6 months	Children from rural areas in a relatively poor province	18 to 30 months old	Cluster level	131 villages 592 children	Endline	DEV: 0.242	Sylvia et al. (2020)
China (2016 to 2017)	Psychosocial stimulation	Weekly individual training in a parenting center by local trainers	12 months	Children from rural areas in a relatively poor province	6 to 24 months old	Cluster level	100 villages 1,720 children	Endline	COG: 0.112	Sylvia et al. (2022)

All impacts but Grantham-McGregor et al.'s (1997) are expressed in standard deviations of the control group. Grantham-McGregor et al.'s (1997) impacts are expressed in standard deviations of the whole sample. The impacts reported all have p -values < 0.1. Some but not all studies adjust p -values for multiple testing.

Abbreviations: COG, cognitive development; COOP, cooperation with test procedures; DEV, cognitive, motor, and socioemotional development; EL, expressive language; FM, fine motor development; H&E, hand-eye coordination; LANG, language; LAZ, length-for-age z -score; LM, locomotor development; M, motor development; NA, not available yet; PRF, behavior performance; RCT, randomized controlled trial; RL, receptive language; SE, socioemotional development; VOC, vocalization; WAZ, weight-for-age z -score.

stimulation improved cognitive development by 0.26 SD (RW p -value 0.002)²⁹ and receptive language by 0.22 SD (RW p -value 0.032) (Attanasio et al. 2014). Similarly sized impacts on cognition were obtained in the home-based Chinese adaptation (Sylvia et al. 2020). In the Cuttack study, the program led to a 0.36 SD (RW p -value 0.016) increase in cognition, 0.26 SD (RW p -value 0.058) increase in receptive language, and 0.21 SD (RW p -value 0.079) increase in expressive language (Andrew et al. 2019).³⁰ In contrast, the original JHV evaluation found that the program led to a 0.9 SD improvement in cognition and 0.6 SD improvement in language (all p -values < 0.01) (Grantham-McGregor et al. 1991). Even greater impacts were found in the Bangladeshi adaptations of Reach Up (e.g., Hamadani et al. 2006, Nahar et al. 2012). In contrast, in the center-based adaptation of the program in China, no significant impacts on infant skills were detected at the end of the intervention (Sylvia et al. 2022).

This observation motivates two comments. First, comparing intervention impacts across studies is inherently challenging (Bond & Lang 2013). In the case of the Reach Up intervention, most studies used the same development measures as primary outcomes. Although this obviously facilitates a comparison of impacts across studies, it does not remove all the challenges associated with comparing impacts across settings. Indeed, two interventions may generate the same impact, but the practice of standardizing impacts with respect to the SD of the control group will make the same impact look a lot bigger in a very homogeneous population (with a small SD) than in a more heterogeneous population (with a larger SD). As proposed by Cunha & Heckman (2008), Cunha et al. (2010), and Bond & Lang (2013), a solution to this problem is to anchor the impact of the intervention to a long-term outcome, such as years of education or wages, which can be measured in more meaningful units. Often, however, this anchoring is hard to achieve due to the scarcity of long-term follow-ups from these interventions.³¹ Another way to achieve a more meaningful comparison of interventions is to measure impacts in terms of the gap in development between well-defined groups in the population (e.g., bottom and top quartiles of the income distribution). This approach requires access to external and, ideally, nationally representative data sets containing the developmental measures used as primary outcomes. These are strong data requirements, but weaker than those necessary for anchoring impacts to long-term outcomes.

Putting aside those methodological considerations for the moment, the fact that impacts of the same intervention were greater in the Jamaican and Bangladeshi contexts than in the other studies could suggest that the program is truly more effective in some populations than in others. As mentioned above, children in the Jamaican and Bangladeshi studies were only eligible to the program if they were undernourished or severely undernourished, as measured by the degree of stunting. The interventions in Colombia and India have generally targeted the poor but not necessarily children suffering from long-term malnutrition. For example, in Colombia the target groups were the beneficiaries of the Familias en Acción conditional cash transfer program. There was little or no stunting in that population, although the children's body mass index was high. The Cuttack intervention targeted slum-dwellers, whereas the rural Odisha intervention targeted village dwellers without any further screening. In both cases, the stunting rate was about 30%.

²⁹RW refers to the p -values adjusted for multiple testing using Romano & Wolf's (2005, 2016) stepdown procedure.

³⁰We report results when the (adjusted) p -value is less than 0.1. If the study does not adjust the p -value for multiple testing, we use the p -value reported in the paper.

³¹One could use an auxiliary data set to perform this anchoring. This would still require the existence of longitudinal panels containing the same developmental measures as those used to measure the impact of the intervention. These data sets are extremely rare in developing countries or anywhere else.

In the production function framework we have discussed in Sections 2 and 3, this would suggest that the marginal productivity of the inputs provided by the intervention (mostly, increased stimulation) may depend negatively on the health endowment of children at baseline. In other words, stimulation may be a substitute for baseline health and cognition.³² And indeed, in the Cuttack study, the impacts were considerably larger (and in line with those found in the JHV and Bangladeshi studies, at about 0.8 SD for cognition) for children that were stunted at baseline (Andrew et al. 2019). From a policy perspective, this would suggest the importance of targeting these interventions to the lowest levels of development, and/or of adapting their content so that even children who start the program with higher levels of development receive appropriate stimulation from it.

It is interesting to consider the heterogeneity of average impacts found not only across studies but also within studies. For example, in the Cuttack study, the intervention was twice as effective in increasing cognitive development for those whose mother had higher levels of education (0.38 SD with RW p -value of 0.04 versus 0.19 SD with RW p -value of 0.31), and the entire impact on receptive language was driven by this group as well (0.37 SD with RW p -value of 0.035) (Andrew et al. 2019). At first, this complementarity between the program and maternal education may seem to contradict the fact that the program may be more effective for children with the lowest levels of development at baseline (since those are also likely to have less educated mothers). But it could also suggest that, in this context, more educated mothers are more likely than less educated mothers to adjust their parenting practices in response to the program (either because they believe in the importance of these practices more or because they have more time and/or less stress in their environment to implement them). While this hypothesis (and, more broadly, the issue of impact heterogeneity across and within interventions) warrants much further investigation, these findings suggest that any production function used to interpret intervention impacts must be flexible enough to allow for complex interactions between the inputs provided by the intervention and the baseline characteristics of children and their primary carers.

5.3. Mechanisms

Making sense of these patterns requires an understanding of the mechanisms through which this type of interventions generates impacts, and the economic framework set out earlier can be helpful in this regard. Within that framework, this type of intervention can be conceptualized in different (though not mutually exclusive) ways. First, it could be modeled as a transfer in kind. Standard economic reasoning would imply that some parental activities and/or expenditures could be crowded out as a result. For example, more time and resources could be shifted to other children in the family unit and/or to the parents themselves. Second, this intervention could be modeled as a shift in parental beliefs about the value of investments and the best manner to achieve them (such as implementing the various activities promoted by the program). If this were the case, the model would predict that the program would lead parents to increase their investments. Third, the program could also be modeled as an intervention that shifts other inputs (outside of parental investments), especially maternal mental health. These effects could arise because of the regular contacts between the mother and the home visitor and/or because the mother feels increased confidence in her parenting. If maternal mental health and parental investments are complementary, the program could deliver even larger impacts by working through those two channels.

³²These results echo the findings of Bitler et al. (2014) in the context of the Head Start study, who find that the impacts of the program were stronger for those children at the bottom of the developmental distribution.

The impacts of these interventions on parental investments, on the quality of the home-learning environment, and on maternal mental health are helpful to suggest which mechanisms are likely to operate. In all studies in which these data were collected, the psychosocial stimulation program was found to increase significantly the quality of the home-learning environment and child-rearing practices.³³ This includes the parenting center-based Chinese intervention, where no significant effects on infant skills were detected (Sylvia et al. 2022). This is important, as it suggests that the intervention does not crowd out parental investments. The Cuttack study and one of the Bangladeshi studies (Hamadani et al. 2019) also found evidence of an improvement in maternal mental health, though this finding was not replicated in the other cases. In the Cuttack study, maternal depressive symptoms, as measured by a shortened version of the Center for Epidemiological Studies-Depression (CES-D) scale, decreased by 0.22 SD (p -value 0.04). In the Bangladeshi study of Hamadani et al. (2019), they decreased by 0.3 SD (p -value 0.05).

However, this evidence is not sufficient to show that the effect of the intervention only operated by shifting parental behaviors and investments. To investigate this, one needs to perform the kind of mediation analysis Attanasio et al. (2020c) conduct for the Colombian study, which also allows for confounding factors by accounting for the endogeneity of investments. This exercise requires estimating the parameters of the production function (or a set of production functions), such as

$$\log Q_{i,t} = \gamma_0 \log Q_{i,t-1} + \gamma_1 T_i + \gamma_2' \log \mathbf{I}_{i,t} + \gamma_3' \mathbf{x}_{i,t} + v_{i,t}, \quad 6.$$

where the output $Q_{i,t}$ denotes cognitive development, and inputs include baseline levels of child development, $Q_{i,t-1}$ (to measure self-productivity and cross-productivity in the case of a multi-dimensional output); the intervention, T_i ; inputs that could be shifted by the intervention, $\mathbf{I}_{i,t}$ (e.g., parental investments and maternal mental health); and other inputs, $\mathbf{x}_{i,t}$, which may be less likely to be shifted by the intervention but could nevertheless be important in the process of child development (e.g., maternal education or the number of siblings living in the household). For simplicity, we assume one endogenous mediating factor, $\mathbf{I}_{i,t}$, though the model could be extended to accommodate several ones and indeed was extended to consider material investments and time investments by Attanasio et al. (2020c). In the framework, the input T_i can be thought of as the direct in-kind transfer provided by the program (i.e., the stimulation provided to the child by the home visitor during the weekly visit).

The question asked is whether the intervention affects $Q_{i,t}$ both directly ($\gamma_1 \neq 0$) and indirectly by shifting investments ($\gamma_2 \neq 0$), or, alternatively, whether only one channel matters (say, the increase in investment). The fundamental difficulty with identifying the mediation channels is the classic economics problem of endogeneity, expressed here by a correlation between $\mathbf{I}_{i,t}$ and $v_{i,t}$. For example, suppose parents compensated for negative shocks to their child's cognition (i.e., a negative realization of $v_{i,t}$) by increasing investments. Estimating Equation 6 by ordinary least squares, thereby ignoring such a phenomenon, would lead to underestimating the effects of investments (i.e., a downward bias in γ_2) and potentially overestimating the direct effect of the intervention.

To solve this standard identification problem, we can either assume that the errors are not correlated, which is the classic mediation analysis used by Heckman et al. (2013), or we need an instrument, denoted $z_{i,t}$, which can be plausibly assumed not to have a direct effect on cognition. The latter approach, which is the one followed by Attanasio et al. (2020c), requires specifying a

³³Time investments are usually measured by the number and frequency of activities, such as reading and playing, parents engage in with the child; material investments are measured by the number of play materials (books, toys) around the house, excluding the materials left by the home visitor.

first-stage equation, which is a reduced-form version of a parental decision rule for investment in the child of the following form:

$$\log \mathbf{I}_{i,t} = \beta_0 + \beta_1 T_i + \beta_2' \mathbf{x}_{i,t} + \beta_3 z_{i,t} + u_{i,t}. \quad 7.$$

By not relying on the exact optimal decision rule, this semistructural approach avoids imposing the restriction that parents know and understand the process of child development, as reflected here in the production function for cognitive skills. This is particularly important, given that the intervention may be shifting parental beliefs about the productivity of investing in and stimulating children.

In addition to being necessary for the identification of γ_2 , the first-stage equation is of interest because it reveals how resources are allocated to children. In this sense, it informs on the origins of inequality in investments and on how the intervention affects investment decisions. In the case of the Colombian study, the treatment effect on investments (measured by β_1) is strong and positive. A zero or a negative value would have been completely consistent with the behavior of a Becker-type altruistic household, but here the evidence shows that far from crowding out parental resources, the intervention causes parents to invest more. Beyond this result, the estimates of the first-stage equation reveal additional important information about the drivers of investment decisions: Children who score higher in cognition at the ages of 1–2 years (at baseline) receive larger investments, consistent with the idea that these are complementary in the production of cognition. Moreover, holding the child's baseline level of cognitive development constant, mothers with higher cognition themselves invest more in their children. This could reflect better levels of understanding of child development, improved availability of resources, and/or a more stable lifestyle. Finally, the presence of other (older) children reduces investments in the subject child (who is the youngest in most cases). The latter may reflect the usual quality-quantity trade-off.

The estimates of the production function are of key importance for understanding the mechanisms through which these interventions operate: In the context of the Colombia intervention, these imply that the entire impact of the intervention operates by increasing parental investment, which in turn improves cognitive development. The direct effect of the treatment is zero, once investments are controlled for. Another noteworthy finding is that the production function is estimated to be Cobb–Douglas, which implies complementarity between maternal education and parental investments. This could be a reason explaining why the impacts of these interventions are often found to be higher among the children of a mother with higher levels of education: Not only do these mothers invest more in their children, but the marginal productivity of their investments is also higher.

The broad implication of this analysis is that interventions are capable of reversing the effects of poverty, at least to an extent, but the findings raise the deeper question as to why poorer parents are investing less than richer parents, when these investments are not (financially) costly. As alluded to earlier, one hypothesis that is consistent with much of the evidence around these interventions is that they work by shifting parental beliefs over the value of investments. Attanasio et al. (2019) provide direct evidence that, in the case of the Colombian study at least, the intervention did precisely that. They use direct measures of parental beliefs over the productivity of investments for children endowed with different levels of skills. With those data, they show that treated parents believe that investments are more productive than parents in the control group. They estimate a structural model to show that this shift in parental beliefs is enough to explain the intervention impacts on parental investments.

5.4. Sustainability of Impacts over Time

As mentioned earlier, the JHV program achieved remarkably strong impacts into adulthood in both cognition and earnings (Gertler et al. 2014). Among the other evaluations of Reach

Up interventions, the only one that has so far collected data some time after the end of the intervention is the study in Colombia. Two years after the intervention ended, none of the benefits that had been observed immediately were visible any more (Andrew et al. 2018).

The production function framework and the evidence on its empirical features can suggest a number of reasons to explain why the effects of the Colombian intervention would fade out over time. As discussed in Section 3.2.2, in a model where dynamics follow a Markov process, persistence of impacts could occur through either the self-productivity or the cross-productivity channel, and/or because the impacts on inputs occurring during the intervention are sustained over time. If this is an accurate characterization of the process of child development, then fade-out could occur because the intervention's initial impacts were too small (given a particular level of self-productivity) and/or because increases in parental investments were not sustained. In the case of the Colombian intervention, both of these factors were likely to be at play.

The initial impacts of the Colombia trial were much smaller than those of the Jamaica trial, and estimates of self-productivity in the production function for cognition was well below one (Attanasio et al. 2020c). Moreover, when measured 2 years after the end of the intervention, the parental investments of the treated group had returned to the same levels as those of the control group (although they had significantly increased during the intervention). If, indeed, parental investments had increased because of a change in parental beliefs about the productivity of investments (as discussed above), this evidence suggests that such change in beliefs may be too narrow or age-specific to ensure sustained changes in environments over time. This suggests that sustaining impacts in the longer run would require to find ways to preserve parental engagement and possibly continue with further interventions, in preschool and after.³⁴

5.5. Scalability of Parenting Interventions

Although efficacy trials, such as the Jamaican study, show the potential of interventions for mitigating and even reversing the effects of poverty on child development, a key challenge lies in designing the intervention to be scalable. Holding the target population constant, scalability is a matter of resource availability and cost, and there are two crucial parameters that affect the cost of the intervention: the human capital of the personnel delivering the intervention and its duration. There is a practical trade-off between achieving strong benefits from the interventions and reducing the implementation cost to the point that governments will be willing to make the investment.

The first issue to consider is the human resources problem. It would be prohibitively expensive to hire college graduates to act as home visitors. However, this may not be necessary. One idea at the core of the interventions discussed above is to use women drawn from the local community (or men, if the local norms allow them to operate in a household setting). These home visitors would have to be trained in delivering the program as designed and acting as mentors for the mothers. However, an important advantage of recruiting home visitors from within the community is their understanding of the local culture and ability to introduce households to the new practices in a culturally appropriate way. Moreover, if properly chosen among influential and trusted individuals, these individuals may act as role models and help promote the new practices in the entire community. The difficulty, of course, relates to training and supervising the home visitors appropriately to ensure they can effectively deliver the intervention and offer support and encouragement to the mothers.

³⁴This is one of the hypotheses being tested in current work in the Odisha study, where children are randomized at the end of the parenting program to an enhanced preschool program or to the status quo.

An alternative model to individual home visits is the delivery of the intervention in groups. The mothers and their children can attend a group session once a week, and a group facilitator can introduce activities that the group of mother and children practice together. There are several potential advantages to such an approach. First, it comes at a fraction of the cost of delivering home visits. Second, in some contexts, group sessions can enhance the formation of networks of otherwise isolated women, thus reinforcing the adoption and improving the acceptability of new parenting practices (Andrew et al. 2020).

A third possible format is center-based parenting interventions, whereby both individual parenting training sessions and structured group activities can be offered in a central location, which also provides a space for caregivers and children to engage in unstructured play.

To date, there is little systematic evidence about which delivery model (group versus home visits) is most cost-effective and which individuals should be optimally chosen to deliver it (as mentors and supervisors), though the answer to these questions is naturally likely to vary depending on the contexts and cultural norms. One of the few exceptions is provided by Sylvia et al. (2022), who compare the effects of the center-based delivery model with the effects of a home-based intervention previously conducted in the same region of rural China (Sylvia et al. 2020), using the same parenting curriculum and public service system. Another exception is the Odisha study, which implemented group sessions in one treatment arm and home visits in another. Strikingly, Grantham-McGregor et al. (2020) report that after 2 years of intervention, the group sessions were equally as effective as the individual home visits, with approximately 0.3 SD improvement in cognition and language. These findings are quite remarkable, particularly because implementing group sessions costs less than 30% as much as delivering home visits. The compliance and attendance rates were much lower among those assigned to group sessions than among those assigned to individual home visits. This implies that the impact of treatment on the treated, scaled up to account for compliance, is much higher for the group intervention than for the home visits, although the compliers for the two intervention types may be different populations. Moreover, the cost of implementing groups at scale should factor in the extra effort required to attract families that did not attend.

The second key parameter underlying the intervention's implementation cost is the length of the intervention. Among the JHV and the various experimental trials of Reach Up, implementation length varied between 6 and 24 months. To our knowledge, there is no experimental evidence to show how effectiveness depends on the intervention's duration. Though not definitive, some insight can be gained from the Odisha study, where outcome data were collected halfway through the intervention, at 12 months. After 1 year of intervention, children in both the group and the home visiting program arms experienced cognitive gains of about 0.3 SD (with RW *p*-values of 0.018 for group sessions and 0.006 for home visits). Children attending groups also showed a 0.31 SD (RW *p*-value 0.006) significant improvement in language, whereas the home visits showed half that improvement, which is not significant (although the two point estimates are not significantly different from each other). After 2 years of intervention, the group modality remained as effective as the individual home visits, with approximately 0.3 SD (RW *p*-values of 0.007 for group sessions and 0.001 for home visits) improvement in cognition and language. Children in both treatment arms were found to have made strong and highly significant improvements in language in the second year, with now an impact of 0.24 SD (RW *p*-value 0.009) on language for home visits. A surprising finding, therefore, is that no further benefit relative to the control group was achieved for cognition in either intervention nor for language in the group sessions. Given the results on fade-out that we have seen from other experiments, including in the JHV and the Colombia study, the second year may have prevented fade-out, and as such it may be particularly important. However, the lack of further progress with respect to the control group in the second

year is troubling and challenging to explain, although it has been observed in other contexts as well (McKay et al. 1978, Yousafzai et al. 2014, Grantham-McGregor & Smith 2016). What causes this plateau in progress, and how can it be overcome? These are questions that remain unanswered but are of key importance if we are to better understand the process of human capital accumulation and its interactions with poverty and intervention.

5.6. The Production Function and Policy Interventions

The optimal timing and duration of policy interventions depend crucially on the process of child development, which economists describe using production functions. Although we have learned a lot over the recent years, the more we learn, the more questions open up. At stake is the design of coherent interventions to improve investments in children, in particular those from lower-income and, broadly defined, deprived backgrounds, so as to address inequality and the intergenerational transmission of poverty.

From the available evidence, we know that the early impacts of several interventions tend to fade out, although there is some evidence of reemergence at a later time in some cases (Bailey et al. 2017, 2020). An implication of these patterns, and in particular of the reemergence of impacts, is that the first-order Markov assumption, whereby all the past can be captured by the current development level of the child, and which is often used in existing studies, may not be a good representation. Furthermore, the short-run evaluation of interventions may only provide a partial and over-pessimistic picture of their effects.

Many estimates of the effects of parental investments on child development imply that early investments are the most potent. However, this evidence does not account for the effects of schools or peers and the dynamic interactions that parental investment (and its effect on several dimensions of development) has with these subsequent inputs. Investments in later childhood and adolescence are likely to be very important and to interact with early parental investments, but these effects have not been measured appropriately, at least by economists. This implies that we would need sequences of programs that last and complement each other throughout childhood while adapting to the demands of each age.

Finally, existing estimates of the production function covering several childhood periods suggest that developmental measures, such as cognition, are increasingly persistent with age. An implication of this evidence, coupled with the fact that the productivity of parental investment changes as children age, could be the existence of opposing forces defining what the right time to intervene should be or, more accurately, how the intensity and duration of interventions should vary throughout childhood. On the one hand, investments at a very young age seem to be highly effective in the short run. On the other hand, lower self-productivity in this period could imply that the impact of investments more easily fades out. A better understanding of the complex dynamics and interactions among different inputs in the process of child development is necessary for the design of effective policies.

6. CONCLUSION AND DIRECTIONS FOR FUTURE RESEARCH

Human capital research has dominated economics ever since Becker (1964) pushed it to the forefront with his seminal work linking human capital to individual income growth, inequality, and intergenerational mobility. While economists were working out the implications of Becker's theory, including the way individuals and families decide to invest in human capital and the implications of such investments, medical, psychological, and neuroscience researchers were working out how the brain develops and establishing the plasticity of intelligence and the importance of the environment in defining individual outcomes. Through these efforts, they uncovered the

importance of early childhood in defining cognitive and socioemotional development and the way early childhood development interacts with later developmental stages, including adolescence, to form adult skills and capabilities.

To use economic language, these discoveries have taught us a lot about the production process of human capital and its complexities. Perhaps some of their most important lessons relate to the plasticity of the human brain and the link between poverty and early developmental deficits, which could be at the origins of the perpetuation of poverty across generations. The literatures in the fields of child development, neuroscience, and economics are now converging and leading to an important interdisciplinary field dedicated to understanding the interaction of human development with socioeconomic conditions and to designing policies that promote the best possible outcomes for children from all backgrounds. This research and policy agenda involves bringing together the lessons from medical research and neuroscience with our understanding as economists of how families make decisions and react to incentives and constraints when investing in their children's development.

Although much progress has been made in recent years, there are still a number of open challenges and research questions, ranging from identifying the details of the process of human development to a full understanding of the behavior of actors involved. The design of effective policies requires a good understanding of how and when the process of human development changes with age, so as to identify where windows of opportunities for effective interventions might be. Another key element for the development and deployment of effective policies is a good understanding of the behaviors of key actors such as parents or teachers. The accumulating evidence is forcing economists to reevaluate and modify models of human capital investments that assume full information on the development process on the part of key actors. In a context where some of these investments are simple and cheap in nature (such as talking, playing, or reading with children), if we assume full information and a complete understanding of the production function of human capital, it is indeed very hard to understand the huge and widening developmental disparities between the poor and the middle class.

This is where interdisciplinarity offers the strongest support to our understanding: Whereas child development specialists need to know about concepts including crowding out and the resource and time constraints affecting parental behavior, economists are sure to fail in policy design and advice without an understanding of developmental complexities and an appreciation of our limited understanding of the returns to child investments in various childhood stages across the income distribution. Moreover, if we are to understand and possibly remedy the disparities across the income distribution, we need to develop a richer model of household behavior than the one that assumes full information on the part of parents and teachers. And we need more research on the formation of beliefs about the human capital production function and how this process varies with income and wealth.

In this review, we discussed the economics literature on the dynamic production functions of human capital and how these functions have been used to learn directly from the data about the process of human development and the productivity of investments in children at various ages. In doing so, we highlighted some of the important challenges and open questions, including the identification of causal links between inputs and outputs in the process of human capital formation when the former are determined endogenously as the product of individual choices. By using data on actual child development, under ideal circumstances at least, this empirical work is intended to reveal the "true" production function reflecting the productivity of the various inputs at different childhood stages. We also discussed extensively the important, but as of yet unresolved, issue of the dimensions of human capital. Following the lead of Heckman and coauthors, most of the economics literature currently focuses on two dimensions: cognitive and socioemotional skills (often

called noncognitive skills in the economics literature). But is that sufficient? Or is it a product of the limited data at our disposal?

These issues are important in the child development literature and have become central to economics as well, following the increased focus on multidimensional skills in the labor market and the role these play in understanding the effects of automation as well as gender disparities (Bernatzky-Koehli 2021). Of course, from an economic point of view, the question is not just whether there are multiple dimensions to intelligence (a key developmental question), but also whether they can be Hicks-aggregated when studying the labor market and the resulting wages. The ability to aggregate skills into, say, one index transforms what is a complex problem in child development into a much simpler one in terms of its economic implications. Nevertheless, the recent economic literature has shown that the relative price of the various skills has been changing over time together with their relative importance as technology changes, implying that the simplification of aggregation is not available.

The other big question, unresolved as far as we are concerned, relates to how we should model parental decisions to invest in children, and in particular how we should incorporate the role of distorted beliefs. We argue that more should be done to incorporate insights from the child development literature around the appropriate characterization and measurement of parental investments. Most data sets used by economists working in this space have coarse measures of resources and time, which seem inadequate to capture important differences across families of various socioeconomic backgrounds and cultures.

Finally, we discussed policy interventions that seek to change parental behavior toward child-rearing and child investments directly. These have generally been shown to be successful in achieving short-run gains. In some cases, but not all, the gains have been shown to be long-lasting. While generally we understand how to structure such interventions to achieve positive effects, there are still important challenges relating to the scaling up of these interventions and the sustainability of their impacts. First, scalability refers not only to the financial cost of running these interventions but also to the ownership and acceptability of the intervention by the community that is targeted. How should interventions be designed and delivered to take account of this important distinction? Second, we need to identify ways to improve outcomes further and to ensure that these improvements are sustained in the longer run: If scaled-up interventions are not capable of producing the kind of outcomes we have seen in the Jamaica intervention and in a few other places, it may be the case that we need to prolong the intervention period and/or complement the early intervention with other interventions in later periods of childhood.

Third, we need to ensure that the intervention can produce benefits both for the hard-to-reach and for the less disadvantaged who have better prior outcomes. As we have discussed earlier in the article, a crucial distinction between the Jamaican and Bangladeshi interventions on the one hand, and the Colombian and Indian interventions on the other, is the fact that, although the latter generally targeted the poor, they did not necessarily target children suffering from long-term malnutrition. An important question, therefore, is whether the design of this intervention is better suited for the ultrapoor and whether modifications could be performed to obtain benefits for less deprived groups. Despite the urgency to improve development for a relatively broad range of initial deficits, it is unlikely that one size will fit all. How can we adapt successful parenting interventions to obtain benefits for a broader range of baseline abilities and levels of deprivation? This question is key for scaling up in cases in which the heterogeneity of the children is likely to be even higher than it was in the interventions discussed earlier. To our knowledge, little is known about this, though it constitutes an important research priority.

Improving child development among the poor is a key challenge for breaking the cycle of poverty. This will require continued research bringing economists and child development

specialists together, with more fieldwork and ever improved and creative approaches. However, from a growth policy perspective it is important to place this in a broader context. Human capital policies are just one element of a set of policies that can promote growth and indeed poverty. For example, without policies that promote entrepreneurship and capital investment, human capital policies may lead to very little because economic opportunity will be absent even for the better educated and skilled. However, absent human capital policies starting at the very beginning of life, growth is likely to be stunted and inequitable.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Diana Perez Lopez for outstanding research assistance. S.C. gratefully acknowledges financial support from the Economic and Social Research Council (ESRC)–funded Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1) and the European Research Council (ERC-2014-CoG-646917-ROMIA); C.M. and O.A. thank the Cowles Foundation and the Institute for Social and Policy Studies at Yale for financial support. We also thank the National Institutes of Health under project R01 HD 72120 for financial support.

LITERATURE CITED

- Acemoglu D, Autor D. 2011. Skills, tasks and technologies: implications for employment and earnings. In *Handbook of Labor Economics*, Vol. 4B, ed. D Card, O Ashenfelter, pp. 1043–171. Amsterdam: Elsevier
- Adhvaryu A, Nyshadham A. 2016. Endowments at birth and parents' investments in children. *Econ. J.* 126(593):781–820
- Agostinelli F, Saharkhiz M, Wiswall MJ. 2020. *Home and school in the development of children*. NBER Work. Pap. 26037
- Agostinelli F, Sorrenti G. 2018. *Money versus time: family income, maternal labor supply, and child development*. Work. Pap. 2018-017, Hum. Cap. Econ. Oppor. Glob. Work. Group, Univ. Chicago, Chicago
- Agostinelli F, Wiswall M. 2016. *Identification of dynamic latent factor models: the implications of re-normalization in a model of child development*. NBER Work. Pap. 22441
- Agostinelli F, Wiswall M. 2020. *Estimating the technology of children's skill formation*. NBER Work. Pap. 22442
- Aizer A, Cunha F. 2012. *The production of human capital: endowments, investments and fertility*. NBER Work. Pap. 18429
- Alesina A, Hohmann S, Michalopoulos S, Papaioannou E. 2021. Intergenerational mobility in Africa. *Econometrica* 89(1). <https://doi.org/10.3982/ECTA17018>
- Almås I, Armand A, Attanasio O, Carneiro P. 2018. Measuring and changing control: women's empowerment and targeted transfers. *Econ. J.* 128(612):F609–39
- Almås I, Attanasio O, Jarvis P. 2021. *Economics and measurement: new measures to model decision making*. Unpubl. Ms., Stockholm Univ., Stockholm, Swed.
- Almlund M, Duckworth AL, Heckman J, Kautz T. 2011. Personality psychology and economics. In *Handbook of the Economics of Education*, Vol. 4, ed. EA Hanushek, S Machin, L Woessmann, pp. 1–181. Amsterdam: Elsevier
- Almond D, Currie J. 2011. Human capital development before age five. In *Handbook of Labor Economics*, Vol. 4B, ed. O Ashenfelter, D Card, pp. 1315–486. Amsterdam: Elsevier
- Almond D, Mazumder B. 2013. Fetal origins and parental responses. *Annu. Rev. Econ.* 5:37–56
- Andrew A, Attanasio O, Augsburg B, Behrman J, Day M, et al. 2020. *Mothers' social networks and socioeconomic gradients of isolation*. NBER Work. Pap. 28049

- Andrew A, Attanasio O, Augsburg B, Day M, Grantham-McGregor S, et al. 2019. Effects of a scalable home-visiting intervention on child development in slums of urban India: evidence from a randomised controlled trial. *J. Child Psychol. Psychiatry* 61(6):644–52
- Andrew A, Attanasio O, Fitzsimons E, Grantham-McGregor S, Meghir C, Rubio-Codina M. 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 Med.* 15(4):e1002556
- Arcidiacono P, Hotz VJ, Kang S. 2012. Modeling college major choices using elicited measures of expectations and counterfactuals. *J. Econom.* 166(1):3–16
- Attanasio O, Bernal R, Giannola M, Nores M. 2020a. *Investment and the changing dynamics of different dimensions*. NBER Work. Pap. 27812
- Attanasio O, Boneva R, Rauh C. 2020b. Parental beliefs about returns to different types of investments in school children. *J. Hum. Resour.* In press. <http://doi.org/10.3368/jhr.58.2.0719-10299R1>
- Attanasio O, Cattán S, Fitzsimons E, Meghir C, Rubio-Codina M. 2020c. Estimating the production function for human capital: results from a randomized controlled trial in Colombia. *Am. Econ. Rev.* 110(1):48–85
- Attanasio O, Cunha F, Jervis P. 2019. *Subjective parental beliefs: their measurement and role*. NBER Work. Pap. 26516
- Attanasio O, de Paula A, Toppeta A. 2020d. *The persistence of socio-emotional skills: life cycle and intergenerational evidence*. NBER Work. Pap. 27823
- Attanasio O, Fernández C, Fitzsimons EOA, Grantham-McGregor SM, Meghir C, Rubio-Codina M. 2014. Using the infrastructure of a conditional cash transfer program to deliver a scalable integrated early child development program in Colombia: cluster randomized controlled trial. *Br. Med. J.* 349:g5785
- Attanasio O, Meghir C, Nix E. 2020e. Human capital development and parental investment in India. *Rev. Econ. Stud.* 87(6):2511–41
- Attanasio O, Meghir C, Nix E, Salvati F. 2017. Human capital growth and poverty: evidence from Ethiopia and Peru. *Rev. Econ. Dyn.* 25:234–59
- Attanasio O, Meghir C, Schady N. 2010. Mexico's conditional cash transfer programme. *Lancet* 375(9719):980
- Bailey DH, Duncan GJ, Cunha F, Foorman BR, Yeager DS. 2020. Persistence and fade-out of educational intervention effects: mechanisms and potential solutions. *Psychol. Sci. Public Interest* 21(2):55–97
- Bailey DH, Duncan GJ, Odgers CL, Yu W. 2017. Persistence and fadeout in the impacts of child and adolescent interventions. *J. Res. Educ. Eff.* 10:7–39
- Baker-Henningham H, López Bóo F. 2010. *Early childhood stimulation interventions in developing countries: a comprehensive literature review*. IZA Discuss. Pap. 5282, Inst. Labor Econ., Bonn, Ger.
- Banerjee A, Duflo E, Goldberg N, Karlan D, Osei R, et al. 2015. A multifaceted program causes lasting progress for the very poor: evidence from six countries. *Science* 348(6236). <http://doi.org/10.1126/science.1260799>
- Bastagli F, Samman E, Both N, Evans M, Sepúlveda M, et al. 2020. *Universal child benefits: policy issues and options*. Tech. Rep., UNICEF-ODI, London/New York
- Becker GS. 1964. *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education*. Chicago: Univ. Chicago Press
- Becker GS, Tomes N. 1986. Human capital and the rise and fall of families. *J. Lab. Econ.* 4(3):S1–39
- Behrman J, Pollak R, Taubman P. 1982. Parental preferences and provision for progeny. *J. Political Econ.* 90(1):52–73
- Bernatzky-Koehli M. 2021. *Labor market behavior and fertility of highly educated women*. Work. Pap., Yale Univ., New Haven, CT
- Berniell L, De la Mata D. 2016. *Skills: measurement and Latin America's current situation and outlook in more skills for work and life: the contributions of families, schools, jobs, and the social environment*. Rep., Dev. Bank Lat. Am., Caracas, Venez.
- Berry J, Dizon-Ross R, Jagnani M. 2020. *Not playing favorites: an experiment on parental fairness preferences*. NBER Work. Pap. 26732
- Bitler M, Hoynes H, Domina T. 2014. *Experimental evidence on distributional impacts of Head Start*. NBER Work. Pap. 20434

- Björklund A, Salvanes KG. 2011. Education and family background: mechanisms and policies. In *Handbook of the Economics of Education*, Vol. 3, ed. EA Hanushek, S Machin, L Woessmann, pp. 201–47. Amsterdam: Elsevier
- Black SE, Devereux PJ. 2011. Recent developments in intergenerational mobility. In *Handbook of Labor Economics*, Vol. 4, ed. D Card, O Ashenfelter, pp. 1487–541. Amsterdam: Elsevier
- Blundell R, Chiappori P, Meghir C. 2005. Collective labor supply with children. *J. Political Econ.* 113(6):1277–306
- Bobba M, Frisnacho V. 2020. *Self-perceptions about academic achievement: evidence from Mexico City*. TSE Work. Pap. 20-1070, Toulouse Sch. Econ., Toulouse, Fr.
- Bond TN, Lang K. 2013. The evolution of the black-white test score gap in grades K-3: the fragility of results. *Rev. Econ. Stat.* 95(5):1468–79
- Boneva T, Rauh C. 2018. Parental beliefs about returns to educational investments—the later the better? *J. Eur. Econ. Assoc.* 16(6):1669–711
- Borghans L, Weel BT, Weinberg BA. 2014. People skills and the labor-market outcomes of underrepresented groups. *ILR Rev.* 67(2):287–334
- Bourguignon F, Chiappori PA. 1992. Collective models of household behavior: an introduction. *Eur. Econ. Rev.* 36(2):355–64
- Boyden J. 2021. *Young Lives: an international study of childhood poverty. Rounds 1–5 constructed files, 2002–2016*. Data Collect., UK Data Serv., Univ. Essex, Essex, UK. 3rd ed. <http://doi.org/10.5255/UKDA-SN-7483-3>
- Butcher KF, Case A. 1994. The effect of sibling sex composition on women's education and earnings. *Q. J. Econ.* 109(3):531–63
- Campbell FA, Ramey CT. 1994. Effects of early intervention on intellectual and academic achievement: a follow-up study of children from low income families. *Child Dev.* 65(2):684–98
- Carneiro P, Cattán S, Nibloe M. 2022. *Early environments and life-cycle profiles of earnings and health: a machine learning approach*. Unpubl. Ms., Univ. Coll. Lond., London
- Carneiro P, García IL, Salvanes KG, Tominey E. 2021. Intergenerational mobility and the timing of parental income. *J. Political Econ.* 129(3). <https://doi.org/10.1086/712443>
- Cattán S, Fitzsimons E, Goodman A, Phimister A, Ploubidis G, Wertz J. 2022. Early childhood inequalities. *IFS Deaton Rev. Inequal.* In press
- Caucutt EM, Lochner L, Mullins J, Park Y. 2020. *Child skill production: accounting for parental and market-based time and goods investments*. NBER Work. Pap. 27838
- Caucutt EM, Lochner L, Park Y. 2017. Correlation, consumption, confusion, or constraints: Why do poor children perform so poorly? *Scand. J. Econ.* 119(1):102–47
- Chaparro J, Sojourner A, Wiswall MJ. 2020. *Early childhood care and cognitive development*. NBER Work. Pap. 26813
- Chevalier A, Gibbons S, Thorpe A, Snell M, Hoskins S. 2009. Students' academic self-perception. *Econ. Educ. Rev.* 28(6):716–27
- Chiappori PA. 1988. Rational household labor supply. *Econometrica* 56(1):63–90
- Chiappori PA, Meghir C. 2015. Intrahousehold inequality. In *Handbook of Income Distribution*, Vol. 2, ed. AB Atkinson, F Bourguignon, pp. 1369–418. Amsterdam: Elsevier
- Cunha F, Elo I, Cullhane J. 2013. *Eliciting maternal expectations about the technology of cognitive skill formation*. NBER Work. Pap. 19144
- Cunha F, Elo I, Cullhane J. 2020. Maternal subjective expectations about the technology of skill formation predict investments in children one year later. *J. Econom.* In press. <https://doi.org/10.1016/j.jeconom.2020.07.044>
- Cunha F, Heckman J. 2008. Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *J. Hum. Resour.* 43(4):738–82
- Cunha F, Heckman J, Lochner L, Masterov D. 2006. Interpreting the evidence on life cycle skill formation. In *Handbook of the Economics of Education*, Vol. 1, ed. E Hanushek, S Machin, L Woessmann, pp. 697–812. Amsterdam: Elsevier
- Cunha F, Heckman J, Schennach S. 2010. Estimating the technology of cognitive and non-cognitive skill formation. *Econometrica* 78(3):883–931

- Dahl GB, Lochner L. 2012. The impact of family income on child achievement: evidence from the Earned Income Tax Credit. *Am. Econ. Rev.* 102(5):1927–56
- De Walque D, Fernald L, Gertler P, Hidrobo M. 2017. Cash transfers and child and adolescent development. In *Child and Adolescent Health and Development*, ed. DAP Bundy, N de Silva, S Horton, DT Jamison, GC Patton, pp. 325–42. Washington, DC: Int. Bank Reconstr. Dev./World Bank
- Del Bono E, Kinsler J, Pavan R. 2020. *Skill formation and the trouble with child non-cognitive skill measures*. IZA Discuss. Pap. 13713, Inst. Labor Econ., Bonn, Ger.
- DelBoca D, Flinn C, Wiswall M. 2014. Household choices and child development. *Rev. Econ. Stud.* 81:137–85
- Deming DJ. 2017. The growing importance of social skills in the labor market. *Q. J. Econ.* 132(4):1593–640
- Dizon-Ross R. 2019. Parents' beliefs about their children's academic ability: implications for educational investments. *Am. Econ. Rev.* 109(8):2728–65
- Doepke M, Sorrenti G, Zilibotti F. 2019. The economics of parenting. *Annu. Rev. Econ.* 11:55–84
- Doepke M, Zilibotti F. 2017. Parenting with style: altruism and paternalism in intergenerational preference transmission. *Econometrica* 85(5):1331–71
- Doepke M, Zilibotti F. 2019. *Love, Money, and Parenting*. Princeton, NJ: Princeton Univ. Press
- Doss C. 2006. The effects of intrahousehold property ownership on expenditure patterns in Ghana. *J. Afr. Econ.* 15(1):149–80
- Dotti Sani GM, Treas J. 2016. Educational gradients in parents' child-care time across countries, 1965–2012. *J. Marriage Fam.* 78(4):1083–96
- Elango S, García J, Heckman J, Hojman A. 2016. Early childhood education. In *Economics of Means-Tested Transfer Programs in the United States*, Vol. 2, ed. RA Moffitt, pp. 235–98. Chicago: Univ. Chicago Press
- Engle PL, Fernald LC, Alderman H, Behrman J, O'Gara C, et al. 2011. Strategies for reducing inequalities and improving developmental outcomes for young children in low-income and middle-income countries. *Lancet* 378(9799):1339–53
- Fernald L, Gertler P, Neufeld L. 2008. Role of cash in conditional cash transfer programmes for child health, growth, and development: an analysis of Mexico's Oportunidades. *Lancet* 371:828–37
- Fiszbein A, Schady N, Ferreira F, Grosh M, Keleher N, et al. 2009. *Conditional Cash Transfers: Reducing Present and Future Poverty*. Washington, DC: World Bank
- Gertler P, Heckman J, Pinto R, Zanolini A, Vermeerch C, et al. 2014. Labor market returns to an early childhood stimulation intervention in Jamaica. *Science* 344(6187):998–1001
- Giannola M. 2021. *Parental investments and intra-household inequality in child human capital: theory, measurement, and evidence from a lab-in-the-field experiment*. Job Mark. Pap., Univ. Coll. Lond., London
- Grantham-McGregor S, Adya A, Attanasio O, Augsburg B, Behrman J, et al. 2020. Group sessions or home visits for early childhood development in India: a cluster RCT. *Pediatrics* 146:e2020002725
- Grantham-McGregor S, Cheung YB, Cueto S, Glewwe P, Richter L, et al. 2007. Developmental potential in the first 5 years for children in developing countries. *Lancet* 369(9555):60–70
- Grantham-McGregor S, Powell C, Walker S, Himes J. 1991. Nutritional supplementation, psychosocial stimulation, and mental development of stunted children: the Jamaican Study. *Lancet* 338(758):1–5
- Grantham-McGregor S, Smith JA. 2016. Extending the Jamaican early childhood development intervention. *J. Appl. Res. Child.* 7(2):4
- Grantham-McGregor S, Walker S, Chang S, Powell C. 1997. Effects of early childhood supplementation with and without stimulation on later development in stunted Jamaican children. *Am. J. Clin. Nutr.* 66(2):247–53
- Hamadani JD, Huda SN, Khatun F, Grantham-McGregor SM. 2006. Psychosocial stimulation improves the development of undernourished children in rural Bangladesh. *J. Nutr.* 136(10):2645–52
- Hamadani JD, Mehrin F, Tofail F, Hasan I, Huda S, et al. 2019. Integrating an early childhood development programme into Bangladeshi primary health-care services: an open-label, cluster-randomised controlled trial. *Lancet Glob. Health* 7(3):366–75
- Heckman JJ, Humphries JE, Kautz T. 2014. *The Myth of Achievement Tests: The GED and the Role of Character in American Life*. Chicago: Univ. Chicago Press
- Heckman JJ, Karapakula G. 2019. *The Perry preschoolers at late midlife: a study in design-specific inference*. NBER Work. Pap. 25888

- Heckman JJ, Moon SH, Pinto R, Savelyev P, Yavitz A. 2010. Analyzing social experiments as implemented: a reexamination of the evidence from the Highscope Perry Preschool Program. *Quant. Econ.* 1(1):1–46
- Heckman JJ, Pinto R, Savelyev P. 2013. Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *Am. Econ. Rev.* 103(6):2052–86
- Heckman JJ, Rubinstein Y. 2001. The importance of noncognitive skills: lessons from the GED testing program. *Am. Econ. Rev.* 91(2):145–49
- Heckman JJ, Stixrud J, Urzua S. 2006. The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *J. Labor Econ.* 24(3):411–82
- Helmerts C, Patnam M. 2011. The formation and evolution of childhood skill acquisition: evidence from India. *J. Dev. Econ.* 95(2):252–66
- Hoddinott J, Haddad L. 1995. Does female income share influence household expenditures? Evidence from Cote d'Ivoire. *Oxf. Bull. Econ. Stat.* 57(1):77–96
- Katz LF, Murphy KM. 1992. Changes in relative wages, 1963–1987: supply and demand factors. *Q. J. Econ.* 107(1):35–78
- Kinsler J, Pavan R. 2021. Local distortions in parental beliefs over child skill. *J. Political Econ.* 129(1):81–100
- Kohn ML. 1963. Social class and parent-child relationships: an interpretation. *Am. J. Sociol.* 68(4):471–80
- Lancet*. 2007. Early child development in developing countries. *Lancet*, Jan. 7. <https://www.thelancet.com/series/early-child-development-in-developing-countries-2007>
- Lareau A. 2003. *Unequal Childhoods: Class, Race, and Family Life*. Berkeley: Univ. Calif. Press
- Leibowitz A. 1974. Home investments in children. *J. Political Econ.* 82(2, Part 2):S111–31
- Lundberg S. 2017. Noncognitive skills as human capital. In *Education, Skills, and Technical Change: Implications for Future US GDP Growth*, ed. CR Hulten, VA Ramey, pp. 219–43. Chicago: Univ. Chicago Press
- Lundberg SJ, Pollak RA, Wales TJ. 1997. Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. *J. Hum. Resour.* 32(3):463–80
- Lynd RS, Lynd HM. 1929. *Middletown: A Study in Contemporary American Culture*. New York: Harcourt, Brace & Co.
- Lynd RS, Lynd HM. 1937. *Middletown in Transition: A Study in Cultural Conflicts*. New York: Harcourt, Brace & Co.
- Macours K, Schady N, Vakis R. 2012. Cash transfers, behavioral changes, and cognitive development in early childhood: evidence from a randomized experiment. *Am. Econ. J. Appl. Econ.* 4(2):247–73
- Macours K, Vakis R. 2010. Seasonal migration and early childhood development. *World Dev.* 38(6):857–69
- Maluccio JA, Flores R. 2005. *Impact evaluation of a conditional cash transfer program: the Nicaraguan Red de Protección Social*. Res. Rep. 141, Int. Food Policy Res. Inst., Washington, DC
- Mani A, Mullainathan S, Shafrir E, Zhao J. 2013. Poverty impedes cognitive function. *Science* 341(6149):976–80
- McKay H, Sinisterra L, McKay A, Gomez H, Lloreda P. 1978. Improving cognitive ability in chronically deprived children. *Science* 200(4339):270–78
- Molina-Millan T, Macours K, Maluccio JA, Tejerina L. 2020. Experimental long-term effects of early-childhood and school-age exposure to a conditional cash transfer program. *J. Dev. Econ.* 143:102385
- Moroni G, Nicoletti C, Tominey E. 2019. *Child socio-emotional skills: the role of parental inputs*. IZA Discuss. Pap. 12432, Inst. Labor Econ., Bonn, Ger.
- Nahar B, Hossain MI, Hamadani J, Ahmed T, Huda S, et al. 2012. Effects of a community-based approach of food and psychosocial stimulation on growth and development of severely malnourished children in Bangladesh: a randomised trial. *Eur. J. Clin. Nutr.* 66(6):701–9
- Olds DL, Holmberg JR, Donelan-McCall N, Luckey DW, Knudtson MD, Robinson J. 2014. Effects of home visits by paraprofessionals and by nurses on children: follow-up of a randomized trial at ages 6 and 9 years. *JAMA Pediatr.* 168(2):114–21
- Olds DL, Kitzman H, Anson B, Smith J, Knudtson M, et al. 2019. Prenatal and infancy nurse home visiting effects on mothers: 18-year follow-up of a randomized trial. *Pediatrics* 144(6):e20183889
- Olds DL, Kitzman H, Cole R, Robinson J, Sidora-Arcoleo K, et al. 2004. Effects of nurse home-visiting on maternal life course and child development: age 6 follow-up results of a randomized trial. *Pediatrics* 114(6):1550–59
- Papageorge NW, Ronda V, Zheng Y. 2019. *The economic value of breaking bad: misbehavior, schooling and the labor market*. NBER Work. Pap. 25602

- Paxson C, Schady N. 2010. Does money matter? The effects of cash transfers on child development in rural Ecuador. *Econ. Dev. Cult. Change* 59(1):187–229
- Putnam RD. 2015. *Our Kids: The American Dream in Crisis*. New York: Simon & Schuster
- Richter L, Black M, Britto P, Daelmans B, Desmond C, et al. 2019. Early childhood development: an imperative for action and measurement at scale. *Br. Med. J. Glob. Health* 4(Suppl. 4):e001302
- Romano JP, Wolf M. 2005. Stepwise multiple testing as formalized data snooping. *Econometrica* 73(4):1237–82
- Romano JP, Wolf M. 2016. Efficient computation of adjusted p -values for resampling-based stepdown multiple testing. *Stat. Probab. Lett.* 113:38–40
- Schady N, Rosero J. 2008. Are cash transfers made to women spent like other sources of income? *Econ. Lett.* 101(3):246–48
- Stinebrickner R, Stinebrickner TR. 2014. A major in science? Initial beliefs and final outcomes for college major and dropout. *Rev. Econ. Stud.* 81(1):426–72
- Stinebrickner TR, Stinebrickner R. 2012. Learning about academic ability and the college dropout decision. *J. Labor Econ.* 30(4):707–48
- Super CM, Herrera MG, Mora JO. 1990. Long-term effects of food supplementation and psychosocial intervention on the physical growth of Colombian infants at risk of malnutrition. *Child Dev.* 61(1):29–49
- Sylvia S, Luo R, Zhong J, Dill SE, Medina A, Rozelle S. 2022. Passive versus active service delivery: comparing the effects of two parenting interventions on early cognitive development in rural China. *World Dev.* 149:105686
- Sylvia S, Warrinnier N, Luo R, Yue A, Attanasio O, et al. 2020. From quantity to quality: delivering a home-based parenting intervention through China's family planning cadres. *Econ. J.* 131(635):1365–400
- Thomas D. 1994. Like father like son, like mother like daughter: parental resources and child height. *J. Hum. Resour.* 29(4):950–88
- Todd PE, Wolpin KI. 2003. On the specification and estimation of the production function for cognitive achievement. *Econ. J.* 113(485):F3–33
- Tofail F, Hamadani JD, Mehrin F, Ridout DA, Huda SN, Grantham-McGregor SM. 2013. Psychosocial stimulation benefits development in nonanemic children but not in anemic, iron-deficient children. *J. Nutr.* 143(6):885–93
- Wang F, Puentes E, Behrman J, Cunha F. 2022. You are what your parents expect: height and local reference points. *J. Econom.* In press. <https://doi.org/10.1016/j.jeconom.2021.09.020>
- Ward-Batts J. 2008. Out of the wallet and into the purse: using micro data to test income pooling. *J. Hum. Resour.* 43(2):325–51
- Yi J, Heckman JJ, Zhang J, Conti G. 2015. Early health shocks, intra-household resource allocation and child outcomes. *Econ. J.* 125(588):F347–71
- Yousafzai AK, Rasheed MA, Rizvi A, Armstrong R, Bhutta ZA. 2014. Effect of integrated responsive stimulation and nutrition interventions in the lady health worker programme in Pakistan on child development, growth, and health outcomes: a cluster-randomised factorial effectiveness trial. *Lancet* 384(9950):1282–93