

The Impact of Child Health and Nutrition on Education in Less Developed Countries (Chapter for the Handbook of Development Economics, Vol. 4)

Paul Glewwe

Department of Applied Economics, University of Minnesota
(pglewwe@apec.umn.edu)

Edward Miguel

Department of Economics, University of California, Berkeley
(emiguel@econ.berkeley.edu)

Keywords: Child health, child nutrition, education, human capital.

JEL codes: I12, I21, O12, O15.

Abstract

Hundreds of millions of children in less developed countries suffer from poor health and nutrition. Children in most less developed countries also complete far fewer years of schooling, and learn less per year of schooling, than do children in developed countries. Recent research has shown that poor health and nutrition among children reduces their time in school and their learning during that time. This implies that programs or policies that increase children's health status could also improve their education outcomes. Given the importance of education for economic development, this link could be a key mechanism to improve the quality of life in less developed countries. Many researchers have attempted to estimate the impact of child health on education outcomes, but there are formidable obstacles to obtaining credible estimates. Data are often scarce, although much less scarce than in previous decades. Even more importantly, there are many possible sources of bias when attempting to estimate relationships between child health and education. This paper provides an overview of what has been learned thus far. Although significant progress has been made, much more research is still needed – especially in estimating the long term impact of child health status on living standards. The chapter first reviews some basic facts about child health and education in less developed countries. It then provides a framework for analyzing the impact of health and nutrition on education, describes estimation problems and potential solutions, and summarizes recent empirical evidence, including both non-experimental and experimental studies. It concludes with suggestions for future research directions.

We would like to thank Gustavo Bobonis, Mark Rosenzweig, John Strauss and T. Paul Schultz for comments on earlier versions of this chapter.

1. Introduction

Many children in less developed countries suffer from poor health and nutrition. The United Nations estimates that one third of preschool age children in less developed countries – a total of 180 million children under age 5 – experience growth stunting relative to international norms (United Nations 2000), while hundreds of millions more suffer from tropical diseases, including malaria and intestinal parasites (WHO 2000). To the extent that poor health and nutrition among children has a negative impact on their education, programs or policies that increase children's health status will also improve their education outcomes. Given the importance of education for economic development (World Bank 2001), this link could be a key mechanism to improve the quality of life for people in less developed countries.

Many researchers have attempted to estimate the impact of child health on education outcomes, but there are formidable obstacles to obtaining credible estimates. Data are often scarce (although they are much less scarce than in previous decades), but even more importantly there are many possible sources of bias when attempting to estimate relationships between child health and education. This paper provides an overview of what has been learned thus far, building on earlier reviews in Behrman (1996) and Glewwe (2005). Although significant progress has been made, much more research is still needed – especially in estimating the long term impact of child health status on living standards.

It is also important to mention at the outset what this chapter does not aim to do. It does not survey the extensive literature on the effects of child birth weight, or other dimensions of intrauterine nutrition and health, on later life outcomes. While many estimation issues are common across these two literatures, the existing birth weight literature focuses almost exclusively on U.S. or other OECD country data, and we thus leave a review of that literature to another forum. For recent work in this area, see Behrman and Rosenzweig (2004) and Almond, Chay, and Lee (2004). Finally, the chapter does not provide a general discussion of education in developing countries; two recent and very thorough general discussions are Glewwe and Kremer (2006) and Orazem, King and Duryea (in this handbook).

The following sections of this chapter first review some basic facts about child health and education in less developed countries, then provide a framework for analyzing the impact of health and nutrition on education, describe estimation problems and potential solutions,

summarize recent empirical evidence, and, finally, make suggestions for future research directions.

2. Some Basic Facts on Health, Nutrition and Education in Less Developed Countries

Children in less developed countries usually have worse health and education outcomes than children in wealthy countries. This reflects the lower incomes of households in these countries, as well as lower quality and less accessible health and education services (relative to wealthy countries). This section sets the stage for the rest of the chapter by presenting some basic patterns on the health, nutrition and education outcomes of children in less developed countries.

2.1 Health and Nutrition

Data on child health and nutrition that are comparable across a wide number of less developed countries are somewhat scarce. The most common data are on nutritional status (based on height and weight) and on mortality. Table 1 provides information on malnutrition (specifically low weight for age) and child mortality. Around the year 2000 (the exact date varies by country), about 27% of children in less developed countries were underweight in the sense that their weight was more than two standard deviations below the median weight of a population of healthy children of the same age. This figure varies widely by region within the less developed world, ranging from 9% in Latin America to 48% in South Asia (note that no figure is available for Sub-Saharan Africa because data are missing for many of the countries in that region). The consensus is that this poor performance in child growth reflects two main factors: inadequate intake of food and repeated episodes of diarrhea.

The mortality figures in Table 1 are sobering. In 1990 10.3% of children born in less developed countries died before they reached the age of five. There is only modest improvement in this figure by 2003; under-five mortality had dropped to 8.7%. There is also wide variation in child mortality across regions. In Latin America the 2003 figure was quite low at only 3.3%, while it was 9.2% in South Asia and a staggering 17.1% in Sub-Saharan Africa, a figure that is almost unchanged since 1990.

There are many different causes of child morbidity and mortality in less developed countries. A recent study by the World Bank (Lopez et al., 2006) presents estimates of the overall “burden of disease” in terms of “disability adjusted life years”, which accounts for both

illnesses and premature mortality, under a set of assumptions about the relative burden of different diseases. This information is given in Table 2, by region and separately for both young (0-4 years) and older (5-14 years) children. Among children age 0-4 in less developed countries, about 15% of total “healthy years of life” are lost either due to mortality or to morbidity. About one half of this overall burden of disease is due to communicable diseases, the most prominent of which are (in descending order of importance) respiratory infections, diarrhea, malaria and measles (the last of which is easily prevented by vaccination). About one fifth of the burden of disease for children age 0 to 4 is from perinatal problems (primarily low birthweight and difficulties during childbirth). Another one fifth reflects non-communicable diseases, which include mental retardation, congenital abnormalities and problems with internal organs. Nutrition problems (other than diarrhea) and injuries (primarily accidents) each account for about 4% of the burden of disease.

The burden of disease for children in the first five years of life shows substantial variation across geographic regions. In East Asia, Latin America and the Middle East, which primarily consist of middle income countries, only about one third (one fourth in the case of Latin America) of the burden of disease is due to communicable diseases, and the percent of healthy years of life lost is about 10%, while in the low income countries of South Asia and Sub-Saharan Africa the burden of disease due to communicable diseases is about one half and three fourths, and the percent of health years lost is 17% and 29%, respectively. AIDS and malaria play only a very small role in all regions except Sub-Saharan Africa, where the combined impact of these two disease accounts for about one fourth of the burden of disease for young children. Finally, three diseases for which children can easily be vaccinated against (pertussis, measles and tetanus) contribute little to the burden of disease in the three regions dominated by middle income countries, while they account for about 8% of the burden of disease in South Asia and about 12% in Sub-Saharan Africa.

The burden of disease is somewhat different for children age 5-14. First, in all regions only about 1 percent of healthy years of life are lost; in each region this percentage is about one twentieth of the respective figures for children age 0-4. This suggests that illnesses of pre-school age children that have permanent effects on children’s mental development may have much stronger effects on education outcomes than illnesses children experience with they are of school age. Second, only about one third (37%) of the burden of disease is due to communicable

diseases, of which respiratory infections and measles are the most prominent. Another one third (30%) is due to non-communicable diseases, and a little less than one third (29%) is due to injuries (again mainly accidents). Again, only about 4% of the burden of disease is due to nutritional problems.

As with younger children, there is wide variation in the burden of disease across geographic regions. In South Asia and Sub-Saharan Africa, which primarily consist of low income countries, communicable diseases contribute more to the overall burden of disease (38% and 56%, respectively). AIDS again plays a big role in Sub-Saharan Africa but a very small role elsewhere, and the role of measles and respiratory infections varies widely.

2.2 Education

Poor health may reduce learning for a variety of reasons, including fewer years enrolled, lower daily attendance, and less efficient learning per day spent in school. This subsection examines recent trends in enrollment and recent data on learning.

School enrollment rates have increased dramatically in almost all less developed countries since 1960 (the earliest year with reliable data), but there is still room for improvement. The most widely available indicator of progress in education is the *gross enrollment rate*, the number of children enrolled in a particular level of education, regardless of age, as a percentage of the population in the age group associated with that level. The age range for primary school is usually 6 to 11 years. In 1960, primary school gross enrollment rates were 65% in low-income countries, 83% in middle-income countries, and over 100% in high-income countries, as seen in Table 3.¹ By 2000, enrollment rates had reached or exceeded 100% in both low and middle income countries, and in all regions except Sub-Saharan Africa, where gross enrollment rates peaked at 80% in 1980 and have declined slightly since, a troubling pattern for the world's poorest region.

Gross enrollment rates above 100% do *not* imply that all school-age children are in school. Both over-reporting and grade repetition can cause reported gross enrollment rates to reach or exceed 100% even when some children never enroll in school. An alternative measure

¹ This classification of countries is defined by per capita income in 1960. *Low-income countries* are those with a per capita income below \$200 per year, *middle-income countries* are those with an income between \$200 and \$450, and *high-income countries* are those with an income greater than \$450. These cutoff points, while arbitrary, yield about the same number of countries in each group.

of progress toward universal primary education is *net enrollment rates*, the number of children enrolled in a particular level of schooling who are of the age associated with that level of schooling, divided by all children of the age associated with that level of schooling. Net enrollment rates can never exceed 100 percent, and they remove the upward bias in gross enrollment rates caused by the enrollment of “overage” children in a given level (due to repetition or delayed enrollment). They do not, however, address overreporting in official data. Table 4 shows that net enrollment rates are much lower than gross enrollment rates for low- and middle-income countries, and net enrollment rates for Sub-Saharan Africa are particularly low at only 56%.

Over the past 40 years, enrollment has increased dramatically at both the primary and secondary levels, as seen in Tables 3 and 5. However, progress in secondary enrollment has slowed in the past two decades. In both low- and middle-income countries the secondary gross enrollment rate increased by about 150% from 1960 to 1980, while the increase from 1980 to 2000 was 59% in low-income countries and about 51% in middle-income countries. Another way to see this is to note that from 1970 to 1980 middle-income countries increased their secondary enrollment ratio from 33% to 51% in only one decade, while low-income countries took 20 years (1980 to 2000) to increase from 34% to 54%. Middle-income countries’ progress slowed down sharply in the 1980s, increasing by only eight percentage points (51% to 59%) in that decade, although the increase was stronger in the 1990s (from 59% to 77%).

Trends in secondary gross enrollment rates from 1960 to 2000 differ substantially by region. The secondary school rates in South Asia, Latin America and the Middle East and North Africa were similar in 1960 (10%, 14%, and 13%, respectively), but by 2000 the rate in Latin America (86%) was much higher than in South Asia (47%) and the Middle East and North Africa (66%). Sub-Saharan Africa’s performance over time has been slower than that of other regions. A final interesting comparison is between Latin America and East Asia. East Asia had a higher secondary enrollment rate than Latin America in 1960 (20% vs. 14%), but the rates in Latin American countries surged in the 1990s, so that the average rate in 2000 was 86%, compared to 67% in East Asia.

In many countries, there are moderate gender disparities in access to education. Slightly more than half, about 56%, of the 113 million school-age children not in school are girls (UNESCO, 2002). As shown in Table 6, primary gross enrollment rate in low-income countries s

are 107% for boys and 98% for girls; this gender gap is somewhat wider at the secondary level, 60% for boys and 47% for girls. In middle-income countries, the primary-school enrollment gap between boys and girls is very small (only 4 percentage points), and in secondary school girls actually have a slightly higher rate than boys. In high-income countries, there is almost no difference in primary enrollment rates, and girls have a slightly higher rate at the secondary level.

Important differences in gender gaps emerge across different regions of the world. In Latin America, East Asia, and Eastern Europe/Former Soviet Union and in the countries in the Organization for Economic Cooperation and Development (OECD), there is almost no gender gap at the primary level, although East Asian countries have a gender gap at the secondary level. In contrast, in Sub-Saharan Africa and Middle East/North Africa, gender gaps are sizable at both the primary and secondary levels. The largest gender gaps at both the primary and the secondary levels are in South Asia.

The figures presented thus far have been on the *quantity* of education; however, the *quality* of education in many less developed countries is low in the sense that children learn much less in school than the curriculum states they should learn. This low quality is not entirely surprising because the rapid expansion of primary and secondary education in less developed countries in recent decades has strained those countries' financial and human resources. Comparisons of education quality across countries require internationally comparable data on academic performance. Two important sources of such data are the Third International Mathematics and Science Study (TIMSS) and Progress in International Reading Literacy Study (PIRLS) projects administered by the International Association for the Evaluation of Educational Achievement (IAEEA).²

The scores of students in grades 7 and 8 on the 1999 TIMSS mathematics test are shown in the first two columns of Table 7. The two developed countries, Japan and the United States, have scores of 579 and 502, respectively. South Korean students scored even higher (587), and Malaysian students also performed very well (519). Scores were generally considerably lower in other less developed countries, ranging from 275 in South Africa to 467 in Thailand. Reading results for grade 4 students in 2001 are shown in the last column of Table 7. All seven of the

² The first and second studies that were precursors to TIMSS were undertaken between 1964 and 1984. The results are not comparable with those of the TIMSS, and very few developing countries were included.

participating less developed countries (Argentina, Belize, Colombia, Iran, Kuwait, Morocco and Turkey) have much lower performance than the three developed countries shown (France, the United Kingdom, and the United States). Note that this pattern occurs despite any sample selection effects caused by lower school enrollment in less developed countries, which if anything is likely to lead more of the poor performers there to miss the exam, thus partially dampening differences across less developed and wealthy countries.

In summary, the health, nutrition and education status of children in less developed countries is much lower than that of their counterparts in developed countries. To formulate policies to improve the status of children in low and middle income countries, a clear understanding of the determinants of health, nutrition and education, and of the impact of health and nutrition on schooling, is needed. The remainder of this chapter assesses what economists and other social scientists have learned in this regard, and provides suggestions for future research.

3. Analytical Framework

This section provides an analytical framework for thinking about the relationships between child health, nutrition and education outcomes. The first subsection presents a simple model of the determinants of children's academic achievement that highlights the role of child health. This is followed by a discussion of the relationships in the model that are of greatest interest to policymakers. For a discussion of these issues that incorporates both child and adult health, but does not focus on the impact child health and nutrition on education outcomes, see the chapters by Mwabu (2007) and Strauss and Thomas (2007) in this handbook.

3.1 A Simple Two-Period Model of Child Health and Schooling Outcomes

To demonstrate the issues that arise when attempting to estimate the impact of child health and nutrition status on schooling outcomes, it is useful to begin with a simple model. Assume that there are two time periods. Time period 1 begins with conception and ends when the child is 5-6 years, the time the child is eligible to enroll in primary school. Time period 2 is the years that the child is of primary school age, say from 6 to 11 years old. (Most research on the impact of health and nutrition on school performance has focused on student performance in primary school.) Although dividing a child's life from conception to age 11 into only two time periods is rather

simplistic, a two-period model illustrates many key issues that must be addressed when attempting to estimate the impact of child health and nutrition on education outcomes.

A useful starting point is a production function for academic skills, as measured by test scores when the child is of primary school-age (time period 2). These skills can be denoted by T_2 . A simple yet very useful specification is the following:

$$T_2 = T_{2,P}(H_1, H_2, EI_1, EI_2, \alpha, SC, YS) \quad (1)$$

where the “P” subscript indicates that this is a production function, H_t is child health in time period t , EI_t is parents’ provision of educational inputs (e.g., school supplies, books, education toys, and—perhaps most importantly—time spent by parents with the child that has pedagogical value) in time period t , α is the child’s innate intelligence (ability), SC is school (and teacher) characteristics, and YS is years of schooling attained in time period 2. All variables have positive impacts on T_2 . For simplicity, school characteristics are assumed not to change over time. Allowing school characteristics (“quality”) to vary over time is somewhat more realistic but would complicate the exposition without making any fundamental contribution to understanding the impact of child health on educational outcomes.

The production function in equation (1) emphasizes the role of child health in determining academic skills. It shows how—holding constant parental education inputs, school characteristics, child ability, and years of schooling—child health status in both time periods could affect learning. This is a *structural* relationship because all of the variables in the production function *directly* affect academic skills, and all the variables with direct effects are included. As will be seen below, *indirect* effects are also possible, but when discussing any effect it is important to distinguish between direct and indirect effects.

If one had accurate data on all the variables in equation (1) one could estimate it using relatively simple methods, such as ordinary least squares, and so obtain unbiased estimates of the direct impacts of all variables, including child health status in both time periods, on child academic skills. To see other relationships that may also be of interest, and to see how these relationships have different data requirements, consider an economic model in which parents maximize the following utility function:

$$U = U(C_1, C_2, H_1, H_2, T_2) \quad (2)$$

where C_t is *parental* consumption of an aggregate consumption good in time period t . Utility is increasing in all variables. For simplicity, this model ignores household utility in later time periods; accounting for later decisions does not change the fundamental insights provided by this model, and in some cases it may be unrealistic to assume that parents make firm plans far into the future. Another simplification is that the amount of leisure consumed is fixed; this implies that actions that use parents' time, such as providing any kind of instruction to their children or taking them to a health care provider, have a price: the wage of the parent whose time is used to carry out those actions.

Utility is maximized subject to a budget constraint, the production function for academic skills shown in equation (1) and two production functions for child health:

$$H_1 = H_{1,p}(C_1^C, M_1, HE_1, \eta) \quad (3)$$

$$H_2 = H_{2,p}(H_1; C_2^C, M_2, HE_2, \eta) \quad (4)$$

where the "P" subscript indicates that this is a production function, C_t^C is the child's consumption of the aggregate consumption good in period t , M_t is health inputs ("medicine" and "medical treatment"), broadly defined, in time period t , HE_t is the local health environment (incidence of infectious diseases, air and water quality, etc.) in time period t , and η is the innate healthiness of the child. All variables in equations (3) and (4) have positive impacts on child health. Assume that both the local health environment and η are beyond the control of the parents. As in the production function for academic skills, these production function relationships include only variables that *directly* affect child health; variables that have only *indirect* effects, such as prices of health inputs or household wealth, are excluded.

The last constraint faced by parents is the intertemporal budget constraint. Let W_0 be the initial wealth of the household, and assume that it can borrow and lend between the two time periods at an interest rate r . The budget constraint (W_0 IMPLICITLY INCLUDES INCOME) is:

$$W_0 = p_{C,1}(C_1 + C_1^C) + p_{C,2}(C_2 + C_2^C)/(1+r) + p_{M,1}M_1 + p_{EI}EI_1 + (p_{M,2}M_2 + p_{EI}EI_2 + p_S YS)/(1+r) \quad (5)$$

Where $p_{C,t}$ is the price of the consumption good in time period t , $p_{M,t}$ is the price of medicine in time period t (which can include the price of travel time, measured in terms of forgone wages, and thus can reflect distance to health facilities), El_t is educational inputs purchased by parents in time period t , p_{EI} is the price of educational inputs (which is assumed constant over time), and p_S is the price of a year of schooling in time period 2.

Optimizing the utility in equation (2) with respect to the constraints in equations (1), (3), (4) and (5) gives the following standard demand functions for the nine endogenous variables that can be purchased in the market:³

$$C_i = C_{i,D}(W_0; r, p_{C,1}, p_{C,2}, p_{M,1}, p_{M,2}, p_{EI}, p_S; HE_1, HE_2, SC, PS; \alpha, \eta, \sigma, \tau) \quad i = 1, 2 \quad (6) \text{ and } (7)$$

$$C_i^C = C_{i,D}^C(W_0; r, p_{C,1}, p_{C,2}, p_{M,1}, p_{M,2}, p_{EI}, p_S; HE_1, HE_2, SC, PS; \alpha, \eta, \sigma, \tau) \quad i = 1, 2 \quad (8) \text{ and } (9)$$

$$M_i = M_{i,D}(W_0; r, p_{C,1}, p_{C,2}, p_{M,1}, p_{M,2}, p_{EI}, p_S; HE_1, HE_2, SC, PS; \alpha, \eta, \sigma, \tau) \quad i = 1, 2 \quad (10) \text{ and } (11)$$

$$El_i = El_{i,D}(W_0; r, p_{C,1}, p_{C,2}, p_{M,1}, p_{M,2}, p_{EI}, p_S; HE_1, HE_2, SC, PS; \alpha, \eta, \sigma, \tau) \quad i = 1, 2 \quad (12) \text{ and } (13)$$

$$YS = YS_D(W_0; r, p_{C,1}, p_{C,2}, p_{M,1}, p_{M,2}, p_{EI}, p_S; HE_1, HE_2, SC, PS; \alpha, \eta, \sigma, \tau) \quad (14)$$

where the “D” subscript indicates that these are (standard) demand functions, PS is parents’ level of schooling, σ is parental tastes for child education, and τ is parental tastes for child health. Parental schooling is added because parent time used to provide instruction to children is likely to be more effective for educated parents, which lowers the effective price (in terms of forgone wages) of providing that educational input. Parental tastes for education and child health reflect variation in the utility function across parents. Note that all of the variables on the right hand side of these demand functions are exogenous; that is, none of them are under the control of the parents.⁴

Another important relationship is the demand for the child’s academic skills. This can be obtained by inserting equations (12), (13) and (14) directly into (1), inserting equations (8) and

³ The term “endogenous” is used here in terms of its meaning in an economic model: endogenous variables are variables that can be influenced by household behavior. Whether these variables are endogenous in an *econometric* sense, that is correlated with the error term in an equation to be estimated, is a separate question, which will be discussed in Section 4.

⁴ Whether these variables are exogenous in the econometric sense of being uncorrelated with the error term in an equation to be estimated is a separate question; this is discussed in Section 4.

(10) into (3), then inserting equations (3), (9) and (11) into (4) and finally inserting equations (3) and (4) into (1):

$$T_2 = T_{2,D}(W_0; r, p_{C,1}, p_{C,2}, p_{M,1}, p_{M,2}, p_{EL}, p_S; HE_1, HE_2, SC, PS; \alpha, \eta, \sigma, \tau) \quad (15)$$

where the “D” subscript indicates that this is a demand equation, and as in the other demand equations all the variables on the right-hand side are exogenous in the sense discussed above.

A final important relationship between child health and educational outcomes is the *conditional demand function* for child academic skills. Suppose that child health in both time periods (i.e. H_1 and H_2) were “fixed” at the utility maximizing levels by “fixing” child consumption and health inputs in both time periods (C_1^C , C_2^C , M_1 and M_2) at the utility maximizing levels (recall that η , HE_1 and HE_2 are exogenously fixed). With the remaining funds (parents are still required to pay for the items that are “fixed”), which can be denoted by W_{CD} , parents will still choose the optimal levels of all the other variables. This gives the following conditional demand function for educational inputs (in both time periods) and years of schooling:

$$EI_1 = EI_{1,CD}(H_1, H_2; W_{CD}, r, p_{C,1}, p_{C,2}, p_{EL}, p_S; SC, PS; \alpha, \eta, \sigma, \tau) = EI_{1,CD}(H_1, H_2; W_{CD}, \omega) \quad (16)$$

$$EI_2 = EI_{2,CD}(H_1, H_2; W_{CD}, r, p_{C,1}, p_{C,2}, p_{EL}, p_S; SC, PS; \alpha, \eta, \sigma, \tau) = EI_{2,CD}(H_1, H_2; W_{CD}, \omega) \quad (17)$$

$$YS = YS_{CD}(H_1, H_2; W_{CD}, r, p_{C,1}, p_{C,2}, p_{EL}, p_S; SC, PS; \alpha, \eta, \sigma, \tau) = YS_{CD}(H_1, H_2; W_{CD}, \omega) \quad (18)$$

where W_{CD} , household “non-health” expenditures, is defined as $W_0 - C_1^C - p_{M,1}M_1 - (p_{C,2}C_2^C - p_{M,2}M_2)/(1+r)$ and the vector ω denotes the vector $\{r, p_{C,1}, p_{C,2}, p_{EL}, p_S; SC, PS; \alpha, \eta, \sigma, \tau\}$.⁵

Inserting these conditional demand functions into the production function for academic skills yields the conditional demand function for those skills:

$$T_2 = T_{2,P}(H_1, H_2, EI_1, EI_2, \alpha, SC, YS) \quad (19)$$

$$= T_{2,P}(H_1, H_2, EI_{1,CD}(H_1, H_2; W_{CD}, \omega), EI_{2,CD}(H_1, H_2; W_{CD}, \omega), \alpha, SC, YS_{CD}(H_1, H_2; W_{CD}, \omega))$$

⁵ The health environment in both time periods does not belong in the conditional demand functions because their only role is to affect health, which already appears in those functions.

$$= T_{2,CD}(H_1, H_2; W_{CD}, \omega, \alpha, SC)$$

This equation shows how, when child health in both time periods is fixed at their utility maximizing levels, small changes in those two variables (holding all exogenous variables constant) affect parents' choice of (demand for) T_2 . Note that these impacts of child health on academic skills are *not* the same as the impacts of child health on academic skills in the production function given in equation (1), because the direct effects measured in equation (1) do not allow for behavioral adjustments to EI_1 , EI_2 and YS , while equation (19) does allow for those adjustments.

Equation (19) can be used to show how small deviations in H_1 or H_2 from their optimal levels will affect the (conditional) demand for academic skills, allowing for behavioral responses by parents. To see how these impacts of child health on academic skills differ from those in equation (1), consider the impact of a small increase in H_2 caused by a “random shock”.⁶

$$\begin{aligned} \frac{\partial T_{2,CD}}{\partial H_2} &= \frac{\partial T_P}{\partial H_2} + \frac{\partial T_P}{\partial EI_1} \frac{\partial EI_{1,CD}}{\partial H_2} + \frac{\partial T_P}{\partial EI_2} \frac{\partial EI_{2,CD}}{\partial H_2} + \frac{\partial T_P}{\partial YS} \frac{\partial YS_{CD}}{\partial H_2} \quad (20) \\ &= \frac{\partial T_P}{\partial H_2} + \frac{\partial T_P}{\partial EI_2} \frac{\partial EI_{2,CD}}{\partial H_2} + \frac{\partial T_P}{\partial YS} \frac{\partial YS_{CD}}{\partial H_2} \end{aligned}$$

There are four distinct impacts of this small exogenous increase in H_2 on academic skills. First, the term $\partial T_P / \partial H_2$ shows that the production function for academic skills will *directly* (and “automatically”) transform this increase in child health in time period 2 into an increase in academic skills. The second term shows how an increase in child health in the second time period, *if this increase is known when decisions are made in the first time period*, will lead to a change in the demand for education inputs in the first period. Yet the assumption that the increase in H_2 is due to a “shock” implies that this increase was not anticipated, so it is

⁶ By “random shock” we mean that the change in H_2 occurs without any change in the endogenous inputs in the health production function for H_2 (H_1 , C_2^C and M_2), which are still held fixed, and without any change in HE_2 or η . Changes in H_2 or H_1 induced by changes in HE_1 , HE_2 , C_1 , C_2 , M_1 or M_2 are discussed below.

impossible for parents to go back in time to alter EI_1 . Thus the second term equals zero because $\partial EI_{1,CD}/\partial H_2$ equals zero.

The third impact works through changes in EI_2 , educational inputs in the second time period, caused by a small exogenous increase in child health in that time period (H_2). The sign of this effect is ambiguous because the sign of $\partial EI_{2,CD}/\partial H_2$ is ambiguous, due to complex income and substitution effects. First, the “automatic” increase in T_2 from the increase in H_2 , via the production function for academic skills, raises “full” income by expanding the household’s consumption possibilities set for C_2 and T_2 ,⁷ which leads to an increase in the demand for T_2 (assuming that T_2 is a normal good). Yet it is possible that the first term in equation (20) raises T_2 by more than this income effect alone warrants, in which case the household would cut back on EI_2 (and YS) to reduce T_2 to the desired level induced by this income effect. Indeed, if H_1 , H_2 and C_1 are weakly separable from C_2 and T_2 in the parents’ utility function, and C_2 and T_2 are normal goods, then parents will shift resources from T_2 to C_2 so that this income effect can be used to increase the consumption of both T_2 and C_2 . Thus the (full) income effect of the exogenous increase in H_2 on $EI_{2,CD}$ (and YS) is likely to be negative, but unless some assumptions are made the impact is ambiguous, even though its impact on T_2 is clearly positive.

In addition to the (full) income effect, there are price effects from an increase in H_2 on T_2 and EI_2 . An increase in H_2 is likely to raise the marginal productivity of EI_2 (and of YS), reducing the shadow price of T_2 and thus increasing its demand, which can be satisfied only by an increase in EI_2 (and in YS). This is an (indirect) own-price effect. There are also (indirect) cross-price effects in that the increase in H_2 will reduce the marginal utility of T_2 if H_2 and T_2 are substitutes in consumption, and thus will reduce the demand for T_2 and consequently the demand for EI_2 and YS . On the other hand, if H_2 and T_2 are complements in consumption, an increase in H_2 will tend to increase the demand for T_2 (and thus for EI_2 and YS). In general, the own-price effects usually outweigh the cross-price effects, so the overall price effect will probably generate an increase in the demand for EI_2 . Yet the (full) income effect on the demand for EI_2 is likely to be negative, therefore the sign of $\partial EI_{2,CD}/\partial H_2$ is also ambiguous, and so the same holds for the sign of the third term in equation (20).

⁷ The assumption that the change in H_2 is a shock that occurs after the time period 1 is over implies that the choice for C_1 cannot be altered, just as the choice for EI_1 could not be altered.

The final term in equation (20) is the impact of H_2 on T_2 via years of schooling. This impact is identical to that of the impact of H_2 via EI_2 . Thus the (full) income effect of an increase in H_2 on YS is ambiguous, although it is likely to be negative. Similarly, the (indirect) own-price effect of an increase in H_2 on YS will be positive, and the sign of the (indirect) cross-price effects will depend on whether H_2 and T_2 are substitutes or compliments in the utility function. While the positive own-price effect is likely to outweigh any negative cross-price effect, so that the overall price effect will be positive, the (full) income effect is likely to be negative, therefore the sign of $\partial YS/\partial H_2$, and thus the overall effect of the fourth term on the demand for T_2 , is ambiguous.

A final comment on the third and fourth terms is that the discussion thus far has assumed that the increase in H_2 occurs before most or all of the decisions regarding EI_2 have been made (or at least before the point is reached that past decisions cannot be reversed). Generally speaking, the later in time period 2 that H_2 is exogenously changed, the less scope there is for changing EI_2 and thus third term will become closer to zero. Yet in contrast to EI_2 , parents will be able to change YS in response to any unexpected increase in H_2 that occurs before the child leaves school, and even if the shock comes after the child leaves school it may be possible for the child to return to school to increase YS . Thus years of schooling is more arguably flexible in its responses to exogenous changes in H_2 than is educational inputs in the second time period.

To summarize, the impact of an exogenous increase in H_2 on the *conditional demand* for T_2 is positive because it raises full income and it tends to reduce the shadow price for T_2 . However, it is unclear whether this impact is larger or smaller than the direct impact through the production function for academic skills (the first term in equation (20)), because the income and price effects on the (conditional) demand for EI_2 and YS work in opposite directions; income effects increase the demand for C_2 and thus lead to a reduction in resources for T_2 (i.e. a reduction in EI_2 and YS), while price effects increase the demand for T_2 . Therefore the net effect of H_2 on T_2 may be either smaller or larger than the direct effect that works through the production function.

Next, consider what happens to the conditional demand for T_2 from an exogenous “shock” that increases H_1 .⁸ Differentiating equation (19) with respect to H_1 yields a somewhat

⁸ As with the change in H_2 discussed above, assume that the change in H_1 is a “shock” that occurs with no change in the purchased or endogenous inputs in the health production function

more complicated expression than that in equation (20), but the overall finding is the same. $\partial T_{CD}/\partial H_1$ is equal to the structural effect, $\partial T_P/\partial H_1 + (\partial T_P/\partial H_2)(\partial H_2/\partial H_1)$, plus three terms that account for behavioral adjustments that alter EI_1 , EI_2 and YS . Income effects suggest that parents will reduce all three of these variables, but price effects create incentives for parents to increase their expenditures on EI_1 , EI_2 and YS . As in the case with an increase in H_2 , the overall impact of an increase in H_1 on T_2 will be positive, but it is unclear whether this increase in the *conditional* demand for T_2 will be greater or smaller than the (aggregate) structural increase that operates via $\partial T_P/\partial H_1 + (\partial T_P/\partial H_2)(\partial H_{2,P}/\partial H_1)$.

It is also instructive to examine how changes in exogenous variables that are likely to increase child health ultimately affect children's academic skills. The net effect of changes in $p_{M,1}$, $p_{M,2}$, HE_1 and HE_2 (prices for health inputs and the overall health environment) on T_2 is obtained directly by differentiating equation (15), yet more can be learned by decomposing these effects to illuminate the pathways by which they take place. Consider first an improvement in the health environment in the second time period, which can be expressed as an increase in HE_2 . Note that such a change has no effect on W_{CD} , since it does not enter the budget constraint.

Substituting (6) and (10) into (3), and (3), (7) and (11) into (4), and then (3) and (4) into (1), and finally (12), (13) and (14) into (1) gives a more detailed *unconditional* demand function for academic skills (T_2). Differentiating this expression with respect to HE_2 gives:

$$\begin{aligned} \frac{\partial T_{2,D}}{\partial HE_2} = & \frac{\partial T_P}{\partial H_2} \left[\frac{\partial H_{2,P}}{\partial HE_2} + \frac{\partial H_{2,P}}{\partial M_2} \frac{\partial M_{2,D}}{\partial HE_2} + \frac{\partial H_{2,P}}{\partial C_2^C} \frac{\partial C_{2,D}^C}{\partial HE_2} \right] \quad (21) \\ & + \frac{\partial H_{2,P}}{\partial H_1} \left(\frac{\partial H_{1,P}}{\partial M_1} \frac{\partial M_{1,D}}{\partial HE_2} + \frac{\partial H_{1,P}}{\partial C_1^C} \frac{\partial C_{1,D}^C}{\partial HE_2} \right) \\ & + \frac{\partial T_{CD}}{\partial H_1} \left[\frac{\partial H_{1,P}}{\partial M_1} \frac{\partial M_{1,D}}{\partial HE_2} + \frac{\partial H_{1,P}}{\partial C_1^C} \frac{\partial C_{1,D}^C}{\partial HE_2} \right] + \frac{\partial T_P}{\partial EI_1} \frac{\partial EI_{1,D}}{\partial HE_2} \end{aligned}$$

for H_1 or H_2 (C_1^C , M_1 , C_2^C and M_2) and without any change in HE_1 , HE_2 or η . Thus this change in H_1 and H_2 does not affect W_{CD} .

$$\begin{aligned}
& + \frac{\partial T_P}{\partial EI_2} \frac{\partial EI_{2,D}}{\partial HE_2} + \frac{\partial T_P}{\partial YS} \frac{\partial YS_D}{\partial HE_2} \\
= & \frac{\partial T_P}{\partial H_2} \left[\frac{\partial H_{2,P}}{\partial HE_2} + \frac{\partial H_{2,P}}{\partial M_2} \frac{\partial M_{2,D}}{\partial HE_2} + \frac{\partial H_{2,P}}{\partial C_2^C} \frac{\partial C_{2,D}^C}{\partial HE_2} \right] + \frac{\partial T_P}{\partial EI_2} \frac{\partial EI_{2,D}}{\partial HE_2} + \frac{\partial T_P}{\partial YS} \frac{\partial YS_D}{\partial HE_2}
\end{aligned}$$

where the last line indicates that a change in the health environment in time period 2 comes “too late” for parents to reverse decisions made in time period 1 (this is relaxed below). Intuitively, a government policy that changes the health environment changes both H_1 and H_2 , but households who are already in time period 2 when the government policy changes cannot change M_1 , C_1^C , C_1 or EI_1 , so for these households there is only a “short-run” effect of the policy change: H_2 changes but not H_1 . A “long-run” effect applies only to households who are still in their first time period when the policy is implemented, or who enter time period 1 after the policy is implemented; for these households both HE_1 and HE_2 change by the same amount, and the long-run impact of that change in the health environment incorporates households’ decisions to change M_1 , C_1^C , C_1 and EI_1 .

Beginning with the short-run effect of an increase in HE_2 on T_2 , the positive structural effect via $(\partial T_P / \partial H_2)(\partial H_{2,P} / \partial HE_2)$ raises the consumption possibilities set in time period 2, so in general there will be a positive income effect. Yet this increase in HE_2 could affect the marginal impacts of M_2 and C_2^C on health in time period 2 (H_2). For example, improvements in sanitation may reduce the incidence of diarrhea and thus render anti-diarrheal medicines less effective in improving health in time period 2, which implies that the demand for those medicines will decrease, so $\partial M_{2,D} / \partial HE_2 < 0$. On the other hand, in some settings improvements in sanitation could make some medicines more effective because such improvements could lower the exposure to infectious diseases and thus increase the duration of improved health from an application of such medicines. Yet even if the overall effect of an increase in HE_2 leads to a decrease in purchases of medicine (i.e. $\partial M_{2,D} / \partial HE_2 < 0$), improvements in sanitation may make addition food more effective in raising child health (for instance, if nutrients are better absorbed by children free of diarrheal disease) and thus will increase the demand for C_2^C , so $\partial C_{2,D}^C / \partial HE_2 > 0$. Overall, the sign of the term $(\partial H_{2,P} / \partial M_2)(\partial M_{2,D} / \partial HE_2) + (\partial H_{2,P} / \partial C_2^C)(\partial C_{2,D}^C / \partial HE_2)$ is ambiguous.

Turning to the last two short run terms, those in the fourth line of equation (21), the overall effect of HE_2 on the demand for E_2 and YS_D is also ambiguous because, as in the case of an exogenous shock to H_2 , the income effects create an incentive to divert resources from producing T_p to increasing of C_2 , while price effects (a reduction in the shadow price of T_2) generate an incentive to increase T_p , so the two variables that can be used to modify T_2 , EI_2 and YS , could either increase or decrease in response to an exogenous increase in child health in the second time period (H_2). In summary, due to multiple ambiguities it is unclear whether the short-run impact of an increase in HE_2 on the demand for T_p will be larger or smaller than the structural impact measured by $(\partial T_p / \partial H_2)(\partial H_{2,p} / \partial HE_2)$.

Now consider the long-run impact of an improvement in the health environment, which amounts to an increase in HE_1 and HE_2 of the same magnitude. With little loss of generality set HE to be the same in both time periods. The long-run impact of an improvement in the health environment is:

$$\begin{aligned}
\frac{\partial T_{2,D}}{\partial HE} &= \frac{\partial T_P}{\partial H_2} \left[\frac{\partial H_{2,P}}{\partial HE} + \frac{\partial H_{2,P}}{\partial M_2} \frac{\partial M_{2,D}}{\partial HE} + \frac{\partial H_{2,P}}{\partial C_2^C} \frac{\partial C_{2,D}^C}{\partial HE} \right. \\
&\quad \left. + \frac{\partial H_{2,P}}{\partial H_1} \left(\frac{\partial H_{1,P}}{\partial HE} + \frac{\partial H_{1,P}}{\partial M_1} \frac{\partial M_{1,D}}{\partial HE} + \frac{\partial H_{1,P}}{\partial C_1^C} \frac{\partial C_{1,D}^C}{\partial HE} \right) \right] \\
\frac{\partial T_{CD}}{\partial H_1} &\left[\frac{\partial H_{1,P}}{\partial M_1} \frac{\partial M_{1,D}}{\partial HE} + \frac{\partial H_{1,P}}{\partial C_1^C} \frac{\partial C_{1,D}^C}{\partial HE} \right] + \frac{\partial T_P}{\partial EI_1} \frac{\partial EI_{1,D}}{\partial HE} \\
&\quad + \frac{\partial T_P}{\partial EI_2} \frac{\partial EI_{2,D}}{\partial HE} + \frac{\partial T_P}{\partial YS} \frac{\partial YS_D}{\partial HE}
\end{aligned} \tag{22}$$

The intuition here is that, in the long run, the change in the health environment also allows the household to adjust child health, parental consumption and education inputs in time period 1. The overall income effect implies that the household will want to increase C_1 , which may take resources away from actions that would otherwise increase T_2 . Moreover, the structural impact

of HE on H_1 , on H_2 and on T_2 (via $\partial T_{2,p}/\partial H_1$, $\partial T_{2,p}/\partial H_2$ and $(\partial T_{2,p}/\partial H_2)(\partial H_{2,p}/\partial H_1)(\partial H_1/\partial HE)$) implies a larger income effect than in the short-run. The issue is whether reallocation of resources to raise C_1 (and perhaps to raise H_1 beyond the direct effect via $\partial H_1/\partial HE$) lowers T_2 (relative to the short-run impact) more than the structural impact via H_1 raises T_2 (relative to the short-run impact). More precisely, the question is whether the net effect of all the terms in the second and third lines of (22) is positive or negative. Using the same reasoning above for the short-run effect, it is difficult to ascertain the sign of $(\partial H_{1,p}/\partial M_1)(\partial M_{1,D}/\partial HE) + (\partial H_{1,p}/\partial C_1^C)(\partial C_{1,D}^C/\partial HE)$, although clearly the impact of $\partial H_1/\partial HE$ will be positive. Turning to the third line, the impact on T_2 of $(\partial H_{1,p}/\partial M_1)(\partial M_{1,D}/\partial HE) + (\partial H_{1,p}/\partial C_1^C)(\partial C_{1,D}^C/\partial HE)$ is equally ambiguous, as is $\partial E_{1,D}/\partial HE$. Thus in the long run as well as the short run, it is unclear whether a change in the health environment will lead to a change in the demand for T_2 that is greater or smaller than the structural impact, which is $(\partial T_p/H_2)(\partial H_{2,p}/\partial HE)$ in the short run and $(\partial T_p/H_2)[(\partial H_{2,p}/\partial HE) + (\partial H_{2,p}/\partial H_1)(\partial H_1/\partial HE)]$ in the long run.

Finally, briefly consider the impact of a government policy to decrease the prices for health inputs, via a reduction in $P_{M,2}$ and $P_{M,1}$. There is no structural effect from this change because prices do not enter directly into the production function for child health. As in the case of a change in the health environment, there are short run and long run effects, the former including only a drop in $P_{M,1}$ while the latter includes a drop in both prices. The overall effect of a drop in either price will be an increase in T_2 , as long as academic skills are a normal good, because a drop in prices increases the effective budget set of the household and also reduces the shadow price of T_2 by reducing the shadow price of H_1 and/or H_2 .

An interesting question is whether a reduction in health input prices in one or both time periods that brings about an improvement in child health that is equal to an improvement brought about by a change in the health environment in one or both time periods. In general, for a given improvement in child health in one or both periods, the increase in child academic skills (T_2) will be higher if the change is induced by a reduction in prices for health inputs. This is the case because the change in the health prices has an income effect that does not occur with a change in the health environment. This can be seen by inspection of the conditional demand relationship in equation (19). The increase in health in one or both time periods has identical effects on child academic skills whether it is brought about by a reduction in prices or an improvement in the health environment. Yet recall that $W_{CD} = W_0 - p_{C,1}C_1^C - p_{M,1}M_1 - (p_{C,2}C_2^C - p_{M,2}M_2)/(1+r)$.

W_{CD} will increase if either price of health inputs decreases, but there is no such effect for a change in the health environment. This income effect will unambiguously increase T_2 , so the impact of a given improvement in child health on child academic skills varies depending on the type of policy that brought it about.

3.2 Relationships of Interest

The previous subsection presented three equations that showed the factors that determine children's academic skills, namely equations (1), (15), and (19). The first is a production function, the second is a standard demand function, and the third is a conditional demand function. Each of these equations depict different processes, and a key question is: Which equation is most useful for making policy decisions? To answer this question, this subsection presents the merits of each of these relationships as guides for policy.

As explained above, equation (1) measures the direct (structural) impact of all variables that have direct impacts, including health status in both time periods, on children's academic skills in time period 2. At first glance, this would appear to be precisely what policymakers would like to know. (Whether this equation can be estimated is a separate question, one that will be discussed in the next section). Yet this relationship *does not necessarily imply* that whenever the government implements a policy that improves a child's health status in one or both of those time periods that the education outcome of that child will increase according to the relationship shown in equation (1). Such discrepancies can arise because changes in child health status may lead parents to change their demand for education inputs and years of schooling, as seen in equations (12), (13), and (14). In particular, the expressions for the change in parents' conditional demand for the child's academic skills due to a change in child health in time period 2 (equation (20)) or in time period 1 (equation (21)), show that the change in those skills, after accounting for behavioral adjustments, could be greater or less than the structural effects obtained by differentiating equation (1) with respect to H_1 or H_2 . Thus while equation (1) is very informative it does *not* necessarily depict what will happen to children's academic skills if a program or policy increases child health in either time period by a certain amount.⁹

⁹ Despite the shortcomings of equation (1) in estimating the actual impact of a change in child health on children's academic skills, it may provide better estimates of the overall welfare benefits of an increase in health because it measures the full impact of that change before parents

The relationship in equation (15) shows how changes in the health environment or in the prices of health inputs (or changes in any other variables in that equation) lead to changes in (the demand for) children's academic skills. Unlike equation (1), this equation accounts for *all* changes in behavior that arise in response to changes in the health environment and in prices of health inputs. For policymakers working in health, equation (15) is precisely what is needed to assess the impact of health *policies*, as opposed to health *status*, on children's academic skills. It measures the overall effect of any health policy or program on children's education outcomes through all potential channels. Thus one need not estimate equation (1) to make policy choices if one has already correctly estimated equation (15). A final caution when using equation (15) to assess the impacts of policy changes is that it is important to distinguish between short-run and long-run effects; these could be quite different if child health in the first time period has much stronger structural effects on educational outcomes than child health in the second period, since the short-run effect includes only the impact on health in the second time period while the long-run effect includes the impacts on both time periods.

Finally, consider the conditional demand relationship in equation (19). It is useful for assessing how "shocks" to health in either time period can affect children's acquisition of academic skills, at least initially before parents adjust child health by modifying the choices of C_t^C (child consumption) and M_t (health inputs). In principle, it can also be used to assess the likely impact of a proposed policy that has a known impact on health but has not yet been implemented; for many new policies, it may be impossible to estimate the reduced form relationship in equation (15) because the data available do not adequately describe the new policy. For example, if clinical trials show how some new type of medicine or health care treatment affects child health in one or both time periods, equation (19) approximates how that medicine or treatment would eventually affect child health in a way that accounts for some, but not all, behavioral choices; the behavioral choices not accounted for would be those associated with child consumption and purchases of health inputs in both time periods (and wealth effects if the policy changes the prices of health inputs). Even if the new policy is implemented, the disadvantage of estimating equation (19) is that it may take 6-8 years before its effect on health

make reallocation decisions. Equations (15) and (19) do not capture the welfare benefits of increased parental consumption that comes from these reallocations, but equation (1) approximates it by an application of the envelope theorem. See Glewwe et al. (2004) for a more detailed explanation.

in the first time period will have had time to affect children's academic skills in time period 2, so using the conditional demand relationship (or, if nothing else is available, the production function in equation (1)) one can approximate the impact of the policy on educational outcomes much more quickly.

In summary, in any analysis of the impact of child health and nutrition on education outcomes, it is important to clarify what relationship one is trying to estimate, and whether the impact is long-run or short-run. Different results in different empirical studies are not necessarily inconsistent; they may be estimates of different relationships, and some may measure long-run impacts while others measure only short-run impacts. In practice, some of these relationships are more difficult to estimate than others; the final choice of what to estimate is determined both by the relationships of interest and by the feasibility of estimating each of those relationships. This brings us to econometric estimation issues, which are reviewed in the next section.

4. Estimation Strategies: Problems and Possible Solutions

While economists know less about education than do education researchers, and certainly know less about health than medical and public health researchers, they have ample experience with, and have rigorously debated, many estimation methods. Economists also know that the methods that can be applied and the relationships that can be estimated depend on the data at hand. This section reviews what can be done to estimate the relationship between child health and education with the three main types of data available: cross-sectional data (data collected from the “real world” at one point in time), panel data (data collected from households or individuals at several points in time, also known as “longitudinal data”), and data from randomized evaluations (data collected from an experiment in which one or more groups is randomly selected to receive a treatment while the non-selected group serves as a control).

4.1 Retrospective Estimates from Cross-sectional Data

The easiest data to collect, and therefore the most common type of data available, are data collected on a large number of children at a single point in time. Such data are often referred to as cross-sectional data, and they usually come from a household survey or a survey of schools. For the purpose of estimating the impact of child health on education outcomes, the minimum requirement for such data is that they contain at least one variable that measures child health and

at least on variable that measures a schooling outcome of interest (often either school enrolment or a score on an academic test). With these two variables alone one can measure correlation at one point in time, but of course correlation does not imply causation.

In fact, to estimate causal relationships that show the impact of child health and nutrition status (or the impact of health policies or programs) on one or more education outcomes -- that is to estimate either equations (1), (15) or (19) -- one needs many more variables. To see why, consider the structural equation (1). To avoid problems of omitted variable bias (this is discussed in more detail below) one needs all of the explanatory variables in that equation that affect education outcomes: health status in both time periods, parental education inputs in both time periods, the child's innate intellectual ability, a large number of school and teacher characteristics, and years of schooling attained.¹⁰ Only the last of these is easy to collect. With cross-sectional data, the only possibility for obtaining child health status and parental education inputs in past years is to ask the children or their parents to recall events from many years ago, which is likely to lead to considerable recall error. Moreover, schools and teachers vary in so many ways that it requires great effort to collect all the relevant data on those variables. Indeed, some school and teacher characteristics are difficult to measure, such as teachers' motivation and principals' managerial ability. Finally, it is not trivial to obtain data on a child's innate ability; even defining that concept is difficult in practice.

Thus, in most cases cross-sectional data will be incomplete in the sense that not all of the variables in equation (1) that determine learning will be in the data set. This is very likely to lead to omitted variable bias in estimates of the impact of child health on education outcomes. For example, suppose that data are available only on the current health status of the primary school student (H_2), not on past health status (H_1). Assume also that the true impact of current health status is small while the impact of past health status is quite large. For example, poor health and nutrition in the first few years of life could have a lasting effect on a child's cognitive development. Because current and past health status are likely to be positively correlated, regressing current test scores on current health status and the non-health variables in equation (1) is likely to produce a positive and statistically significant coefficient on current health status,

¹⁰ The assumption in the model of Section 3 that there are only two time periods in the child's life up to age 11 was imposed solely to simplify the exposition. More realistic models are likely to need more time periods, which implies that health status and parental educational inputs must be measured for three or more time periods over the child's life up to age 11.

overestimating the true impact of current health status. If not interpreted cautiously, this could persuade policymakers to put large resources into programs that attempt to improve the current health of school-age children even though programs that focus on infants and very young children may be much more effective.

Another example of possible omitted variable bias is bias due to endogenous program placement when estimating the impact of a health program as in equation (15), where the health program affects child health by altering the health environment or the price of health inputs. Suppose that one has incomplete data on aspects of the local health environment (HE) pertaining to the natural prevalence of childhood diseases. Governments may attempt to address this problem by implementing a program (which would also be an HE variable) to reduce the prevalence of one or more of those diseases. Assuming that the program works, the intervention will have a negative causal impact on the prevalence of the childhood disease(s), and thus a positive impact on child health and on subsequent education outcomes. But if one observes only the program variable, and not the incidence of childhood diseases the estimate of the impact of that variable on children's education outcomes will be biased downward. Intuitively, if the program is implemented primarily in areas with high disease prevalence, this produces a positive association between the program and the prevalence of that disease and thus a negative association between the program and children's academic performance.

Omitted variable bias can also occur when non-health variables are missing. Suppose that parents of some healthy children understand that their children will do relatively well in school without additional investments, and thus they decide to reduce their efforts, and expenditures, on education inputs. This would lead to underestimation of the impact of child health on education outcomes in the structural equation (1) if the data do not include important components of parents' education inputs.

Another plausible example is that parental tastes for child education and child health are correlated, for example some parents are more "responsible" than others, caring about both the health and the education of their children. These tastes are difficult to observe, which will result in positive correlations between child health and child education that are not directly causal. Stated more crudely, irresponsible parents are likely to have children who are both less healthy and do less well in school than the children of responsible parents, but much (perhaps even most)

of the causality may be from parental tastes (more specifically, the actions those tastes produce) to child education outcomes, not only from the direct impact of poor child health on schooling.

In addition to omitted variable bias, another estimation problem can arise: random measurement errors in the explanatory variables can lead to underestimation of the impact of the poorly measured variables on education outcomes (attenuation bias). If cross-sectional data include any retrospective data on past health status and parental education inputs, these data are likely to be measured with a substantial amount of error and thus estimates based on them are likely to suffer from bias towards zero (if measurement error is classical) or bias in an unknown direction (if measurement is non-classical, which is plausible in the context of retrospective health and education reports). Even current health status and parental education inputs may be measured with error, as could current school quality variables.

The discussion thus far has focused primarily on the structural equation (1), but the same estimation problems apply to the demand relationship in equation (15) and the conditional demand relationship in equation (19). Equation (15) has the advantage that parental education inputs are replaced by variables that are probably easier to observe (and thus to collect data on), such as household wealth, parental education, and prices of health and education inputs. Yet other hard to observe factors also appear, such as parental tastes for child education and health (σ and τ) and the child's innate healthiness (η) and innate ability (α), so omitted variable bias remains a very real problem; indeed, it is not clear whether the potential for such bias is lower in equation (15) than in equation (1). Measurement error is also a potential problem, and it is likely to be serious for variables in equation (15) that are not in equation (1), such as household wealth, prices, and the health environment. On the other hand, one could argue that the impacts of the price of medical care and the health environment variables, and of τ and η , in equation (15) are likely to be small and thus these can be dropped from that equation. This may be correct in some settings but we know of no study that has attempted to test the plausibility of this conjecture.

Finally, the demand equation (15) may be easier to estimate than the conditional demand equation (19) because the child health variables need not be directly observed. Yet equation (19) does not include the health environment variables and prices of health inputs (since it conditions on child health), so the endogenous program placement bias problem can occur in estimates of the demand equation (15) but not in estimates of the conditional demand equation (19).

Moreover, the health environment can vary in dozens if not hundreds of different ways, which could imply major data collection difficulties when the goal is to estimate equation (15).

The standard econometric tool for overcoming bias due to omitted variables (other than collecting data on virtually all variables, which may never be possible) and for removing bias due to random measurement error in the explanatory variables is instrumental variable (IV) estimation. The basic idea is that all unobserved variables and errors in measurement can be considered to be included in the error term (residual) of the regression model, and the bias is due to correlation of the observed variables with that error term. If one can find valid instrumental variables—that is, variables that are: (1) correlated with the observed variables that are likely to have bias problems (the relevance condition); (2) uncorrelated with the error term, and thus uncorrelated with all unobserved variables and any measurement errors (the exogeneity condition); and (3) not already included as explanatory variables in the equation of interest (the exclusion restriction)—one can then obtain unbiased estimates by first regressing the observed endogenous variables on the instruments, and then using the predicted values of these observed variables (instead of their actual values) as regressors in the equation of interest.

While IV estimation works in theory, it is very hard to find plausible instrumental variables for use in cross-sectional estimation of the impact of child health on education. Suppose, for example, that one is trying to estimate the structural relationship between child health and education in equation (1), and there are data on child health in time period 2 but not for the earlier time period. As mentioned above, child health is likely to be positively correlated over time, which will lead to overestimation of the impact of child health in time period 2 on students' current academic skills if simple ordinary least squares (OLS) estimation is used. The IV method requires an instrument that predicts child health in time period 2 but is not correlated with child health in the earlier time period. At first glance, health prices in time period 2 seem to satisfy these criteria, but health prices may change little over time and thus those prices could be highly correlated with health prices, and thus with child health, in the first period. Other examples of problems finding valid instrumental variables will be discussed below. On the other hand, the model in Section 3 provides a theoretical argument for a set of instrumental variables to estimate the conditional demand relationship in equation (19); the health environment (HE_t) and health input price ($P_{M,t}$) variables clearly satisfy the exclusion restriction (they do not affect child academic skills after conditioning on child health) and should satisfy the relevance

condition and thus can be used as instruments for child health (H) in both time periods (although the exogeneity condition must still be examined).

4.2 Retrospective estimates from panel data

Panel data are data collected on the same children for two or more time periods. Researchers interested in the impact of child health on education outcomes have an obvious reason for using such data to estimate equations (1), (15), and (19), which is that all three equations include not only variables from the second time period but also variables from the first time period. As pointed out above, cross-sectional data can include such variables only if they are obtained from respondents' memories, which could often be quite inaccurate (imagine trying to remember your exact height in centimeters when you were ten years old). Panel data need not be based on respondents' likely flawed memories of past events.

There is another potential benefit of panel data, which is that some unobserved variables that do not change over time can be differenced out of the regression and thus need not be measured. Estimates of equations (1), (15) and (19) using cross sectional data can lead to biased estimates because many variables that do not change over time – such as child intellectual ability and innate healthiness, parental tastes for educated and healthy children, and some aspects of school quality – are not observed and could be correlated with observed child health outcomes, leading to omitted variable bias. Similarly, any such variables that are measured with error are likely to lead to attenuation bias. In principle, panel data allows one to difference out these unchanging variables and estimate relationships of interest between the variables that do change over time. In many cases, the variables that change over time may also be relatively easier to observe than fixed characteristics (e.g., innate child healthiness), so the omitted variable bias problem is likely to be reduced.

However, this method has its own limitations. It assumes that the troublesome unobserved variables do not change over time, and that they do not interact with variables that do change over time. If either of these assumptions is untrue, then those variables will remain (and will still be unobserved) in the equation being estimated, leading to bias. Another serious problem is that measurement error in observed explanatory variables could lead to greater attenuation bias in estimates based on differenced equations than in estimates based on the original equation, if the signal to noise ratio is smaller for the differenced variables than for the

variables themselves. There is also the obvious disadvantage that panel data are more expensive to collect because they require collecting data at two or more points in time. Limiting sample attrition in panel data collection is often expensive – and challenging – in practice, since respondents who have moved need to be located and interviewed. Movers are often an interesting and highly selected group, and thus important for drawing valid econometric inference. In particular, sample attrition may lead to biased estimates if tracking success is correlated with the variables of interest in the estimation equation, for instance, if healthier individuals are more (or less) likely to migrate elsewhere for work.

For a detailed discussion of the benefits and limitations of panel data, as well as practical advice for collecting such data in less developed countries, see Glewwe and Jacoby (2000). Further examples of how panel data can be used to estimate the impact of child health on education outcomes are discussed in Section 5.

4.3 Randomized evaluations

In the vast majority of studies, both cross-sectional data and panel data are collected from observational settings, that is, settings in which no attempt is made by the researchers to alter the behavior of the people from whom the data are collected, known as observational settings. Yet the problems of bias raised above are very likely, if not almost certain, when using data collected in this manner.

A very different approach to estimating the impact of policies and programs is a method that has long been used in medical sciences: randomized evaluation. Randomized evaluations randomly divide a population under study into two groups, one of which participates in the program, called the treatment group (or program group), and the other of which does not participate in the program, the control group (or comparison group). In some cases the population is divided into more than two groups, one control group and several treatment groups, each with a different treatment. If the division of the population into these groups is truly random, then the *only* difference between the two groups (other than random variation) is that one participated in the program while the other did not. While randomized studies have long been used in health research, until recently they have been rare in social science research, including economic research. Randomized evaluations provide particularly transparent and credible evidence to policymakers on program impacts, and have the potential to exert

considerable influence on actual policy choices, as argued recently by Kremer (2003) and by Duflo and Kremer (2007) in this handbook.

To see how randomized evaluations can be used to estimate the impact of child health on education outcomes, consider the demand for the child's academic skills, as shown in equation (15). A large sample of households or schools can be randomly divided into two groups, a treatment group that receives the health intervention (which can be characterized formally as a change in one or more of the health input price or health environment variables), and a control group that does not receive the intervention. The differences across these two groups in the variables that characterize the intervention are completely uncorrelated with all of the other explanatory variables because these differences are determined solely by random assignment. Thus the difference in the average education outcomes (T_2) of the two groups must be due to the health intervention, since there are no other systematic differences between the two groups. This same logic applies to subgroups of interest within the general population: one can estimate impacts separately by sex, wealth level, or any other group that can be defined using exogenous variables, or using any endogenous variables that are measured before the intervention is implemented.

While this may appear to be the solution to the econometric problems that stymie attempts to estimate such impacts from cross-sectional or panel data, randomized evaluations also have some limitations. First, they are limited to health interventions that do not violate regulations on human subjects research. In health studies this stricture often is interpreted to mean that anyone who is known to have a treatable health problem cannot be denied access to any treatment that is being made available to others. Second, random assignment to treatment and control groups is often violated in practice, as individuals or households in the control group attempt to switch into the treatment group. Even if researchers exclude from the analysis children who were randomly assigned to the control group but were able to obtain the treatment (e.g., enrolled in a treatment school), such children could affect the impact of the treatment on the children who were randomly selected to receive the treatment (e.g., by increasing class size in the treatment schools). This problem can often be addressed in practice, however, by an application of the approach in Angrist, Imbens, and Rubin (1995), in which assignment to treatment is used as an instrumental variable for actual treatment, under fairly weak assumptions.

Third, as with studies based on panel data, randomized evaluations may suffer from attrition bias, and this could lead to bias if attrition is correlated with a child's treatment status. For instance, if the health intervention makes schooling more attractive, the dropout rate among the treatment group may decline. If the study is based on a sample of schools, weaker students will be less likely to drop out of the treatment schools (and thus typically out of the sample) than weaker students in the control schools, and so over time the impact of the program on student academic skills will be underestimated because the average innate ability of students in the treatment schools gradually drops relative to the average ability in the control schools. This sort of differential attrition need not be fatal for estimation, however. Under certain assumptions, researchers can place bounds on the resulting treatment effect estimates using the nonparametric methods described in Manski (1995) and the trimming method in Lee (2002). Unfortunately, in some cases these bounds may be too wide to be useful in practice.

A final limitation of randomized trials is that they are typically designed to estimate only the demand relationship in equation (15), more specifically they estimate the net effect of changes in one or more of the health input price ($P_{M,t}$) and health environment (HE_t) variables. Even if additional data are collected, they usually cannot be used to estimate the structural (direct) impact of child health status on education in equation (1) because one cannot disaggregate the overall impact of the intervention into the effects that work through the various elements of child health status (H_i) and the effects that operate through parental educational inputs (EI_i) and years of schooling (YS). If data collection is extensive, however, one may be able to combine data from randomized trials with structural modeling to recover estimates of key theoretical parameters of interest (for an example of a related approach in another context, see Todd and Wolpin 2003). An important consequence of this limitation is that there may be a long time lag between the start of the intervention and the evaluation of its impact, as seen in the example above concerning an intervention that occurs in early childhood and thus requires 6–8 years before one can evaluate its impact on learning in primary school.

One strength of randomized evaluations is that they can be conducted with only one round of data collection—that is, by collecting cross-sectional data after the health policy or program has been implemented for the treatment group (and after enough time has passed to allow the intervention to have some effect). Another approach, which may be more statistically efficient is to collect panel data that measure children's education outcomes for the treatment and

the control groups both before and after the intervention has been implemented in the treatment group. This allows researchers to look at *changes* in the outcome variables over time, which in some cases will provide an estimate of the impact of the program that has a smaller standard error.

5. Empirical evidence

This section reviews recent studies that examine the impact of child health and/or nutritional status on education outcomes. This is done for all three estimation methods (using cross-sectional data, using panel data and using randomized evaluations). For each method, the studies examined are among the best analyses done in recent years.

5.1 Retrospective estimates using cross-sectional data

Over the past 20–30 years, many studies have attempted to estimate the impact of child health status on education outcomes using cross-sectional data. Yet, as noted by Behrman (1996), most of these studies, especially the earlier ones, paid little attention to the possible biases that can arise when using cross-sectional data, and Behrman concludes that “because associations in cross-sectional data may substantially over- or underestimate true causal effects, however, much less is known about the subject than is presumed” (page 24).

This subsection examines a paper by Glewwe and Jacoby (1995) that carefully investigates the impact of child nutrition on age of school enrollment and years of completed schooling using cross-sectional data from Ghana. Although the paper did not examine the impact of child nutrition on academic skills, the estimation issues encountered in the paper are virtually identical to those discussed above. Thus this paper is instructive in that it shows what can be done, and what cannot be done, using cross-sectional data. This paper is also typical for this literature in focusing on school enrollment derived from household surveys as the main educational outcome measure; we discuss alternative education data below.

Glewwe and Jacoby investigate delayed enrollment and (ultimate) grade attainment using cross-sectional data on 1,757 Ghanaian children aged 6–15 years in 1988–89. They use child height-for-age as their indicator of child health status; in terms of the model in Section 3, this variable reflects health status in both time periods but is primarily influenced by child health in the first time period (more precisely, in the first two or three years of that time period). As

explained above, one problem with using cross-sectional data is that parental tastes for child health and child education outcomes (τ and σ , respectively) may be positively correlated. Glewwe and Jacoby propose a simple way to avoid such bias: they use only variation *within* families, not across families, to estimate the impact of child health on education outcomes. In particular, there is evidence that child health varies within families, but since parental tastes for child health and education outcomes do not vary within the family, within family correlation of child health and education outcomes should not be caused by any such correlation in parental tastes. A family fixed effects estimation procedure can be used to provide estimates that are based solely on within-family variation in health and education outcomes. This is very similar to the differencing approach for panel data discussed above in subsection 4.2, the only difference being that the differences are not over time for one child but instead are across two children in the same family at the same time. Since the two dependent variables, delayed school enrollment and eventual years of schooling, reflect preferences and optimizing behavior, all the relationships estimated in this paper are conditional demand relationships similar to equation (19) in Section 3, rather than structural estimates of education production functions or unconditional demand relationship (equations (1) and (15), respectively).

Another approach used by Glewwe and Jacoby to avoid biased estimates of the impact of child health on education outcomes is to search for instrumental variables that affect child health status but should have no causal impact on education outcomes after conditioning on (controlling for) child health status. The instrumental variables used are distance to nearby medical facilities and maternal height. However, this method can be used only when analyzing variation across households, since these instruments do not vary across children in the same family, and thus this method complements the household fixed effect approach discussed above. Distance to nearby medical facilities, which can be thought of as one of the health input price variables ($P_{M,i}$), should have an effect on child height, while mother's height reflects the mother's, and thus the child's, innate (genetic) healthiness (η). The key assumption, which follows from the theory of conditional demand relationships, is that the price variable ($P_{M,i}$) and innate genetic healthiness (η) can be removed from the list of exogenous variables for those relationships because they affect child schooling only through their impact on child health status.

Yet both of these approaches are open to reasonable criticisms. The authors admit that the first approach (family fixed effects), has a serious problem: variation in innate child healthiness

(η) or random shocks to health among children within the same family may lead to reallocation of (unobserved) education resources across different children within that family. For example, suppose that parents recognize that their children who are relatively sickly will do worse in school. In response, they may allocate more (unobserved) education resources to that child to compensate for the disadvantage the child has in terms of his or her health. Family fixed effects estimation will not control for this intra-household allocation and, in this case, will tend to underestimate the impact of child health on education outcomes. Alternatively, if families decide to neglect sickly children and allocate most education resources to healthier children, then the impact of child health on education outcomes would be overestimated. In the absence of detailed data on intra-household allocation of resources including parental time, it is impossible to account for this effect. This casts doubt on the main results in Glewwe and Jacoby (1995) and, more generally, illustrates the limitations of cross-sectional analysis.

Turning to the second approach (instrumental variables), consider the conditional demand relationship for years of schooling. (The following line of argument also applies to the conditional demand relationship for delayed school enrollment, which is simply another parental choice made in the second time period). This is equation (18) of Section 3. The assumption that the height of the mother and the distance to the nearest medical facility affect only child health is doubtful. A mother's height is likely to influence the marginal productivity of her labor, which affects household income and could influence unobserved parental time devoted to the educational activities of her children. While the distance to the nearest medical facility may affect schooling only through its impact on child health, it could be correlated with many community characteristics that influence education decisions, such as unobserved components of school quality. Thus both instrumental variables are likely to be correlated with the error term when estimating equation (18).

After explaining the limitations of their empirical work, Glewwe and Jacoby estimate the impact of child health (as measured by height-for-age) on school enrollment and final school attainment. They find strong negative impacts of child health on delayed enrollment using both the instrumental variable and fixed effects estimators, and they find little evidence for alternative explanations for delayed enrollment (credit constraints or rationing of limited spaces in school by child age). However, they find no statistically significant evidence that child health increases school attainment – indeed, the point estimate has an unexpected negative sign, although only

marginally significant – but this may reflect the small sample size, since only about 7% of the children in the sample had finished their schooling at the time of the survey.

The above caveats are not limited to estimates of the determinants of years of schooling or delayed enrollment. In general, all of the above discussion applies with very little modification if one were to use cross-sectional data to estimate the impact of child health and nutrition on children's academic skills as measured by test scores. Clearly, very strong and often untestable assumptions need to be made for inference using cross-sectional data.

5.2 Retrospective estimates using panel data

Three recent studies have used panel data to estimate the impact of child health on education outcomes. The first, by Alderman et al. (2001), uses panel data collected from 1986 to 1991 for about 800 households in rural Pakistan. To avoid biased estimates due to unobserved parental tastes and children's innate ability and healthiness, the paper uses food prices (more precisely, deviations in prices from long-term trends) during time period 1 as instrumental variables for child health status in that time period. Education decisions in the second time period are assumed to be made conditional on all outcomes at the end of time period 1, which reflect not only decisions made in the earlier time period but also various exogenous shocks that occurred after decisions were made in the first time period.

Alderman and his coauthors find that child health, as measured by height-for-age when 5 years old, has a strong positive effect on the probability of being enrolled in school at age 7, especially for girls. This finding is consistent with the Glewwe and Jacoby results from Ghana that better health reduces delayed enrollment, since part of the enrollment impact in Pakistan is likely to operate through reducing delayed enrollment. More generally, the results for the two countries are consistent in the sense that improved child health, as captured by height, appears to have a large positive causal impact on education outcomes.

The Pakistan study has several potential limitations. First, the relationship they estimate is a conditional demand function. Thus their use of food price shocks in the first time period as instrumental variables for health status in that time period is theoretically valid only if they include an initial wealth variable that excludes spending on child health in the first time period (i.e. the appropriate wealth variable in terms of the model in Section 3 is W_{CD} , not W_0). But their household wealth variable, household expenditures averaged over three years, does not exclude

spending on child health. Moreover, food prices in the first time period determine not only child health but also adult food consumption and thus they belong in the conditional demand function even after controlling for child health ($p_{C,1}$ is part of the ω vector in equation (18)). More intuitively, food price shocks can affect household savings in the first time period and thus affect education choices in the second time period. Thus, the use of food price shocks as instruments for health outcomes in equation (1) potentially violates the exclusion restriction, the requirement that the instruments have no effect on years of schooling apart from the effect that operates via lagged health status (H_1). The direction of bias is toward overestimation of the health effects: unusually high prices in the first time period probably not only reduce child health but also reduce savings for education inputs, via an income effect.

Another potential concern is that the paper assumes that household wealth (as proxied by consumption expenditures) is measured without error. Yet it is very likely that at least some measurement error is present, which implies biased estimates of the impacts not only of the consumption variable but also potentially of all other variables. Addressing this issue would require finding a suitable instrumental variable for consumption expenditures.

A second recent paper using panel data is that of Glewwe, Jacoby, and King (2001), which uses panel data from more than 2,000 households in the Philippines. Unlike the Ghana and Pakistan studies, this paper estimates the determinants of academic skills as measured by test scores, not school enrollment, and it attempts to estimate the structural educational production function in equation (1), as opposed to estimating a conditional demand function. By making certain assumptions the authors attempt to get around the problem that the instruments could be correlated with unobserved parental education inputs in the first time period.

The Philippines study, like the Ghana study, is based on sibling differences. As will be seen below, this differencing is useful because it removes family averages of innate academic ability (α) and all school quality variables (virtually all siblings in the sample attended the same primary school) from equation (1). Yet using household fixed effects does not remove bias due to possible differences in innate ability across different children in the same family. In particular, decisions regarding health investments throughout childhood ($M_1 C_1^C$, C_2^C), as well as decisions on educational inputs (EI_1 and EI_2), could be influenced by differences in innate ability among siblings in the same family, which may lead to correlation between early childhood health investments and primary school test scores (T_2) that once again is not causal. The authors argue

that health investments made from conception through 24 months of age cannot be correlated with innate child academic ability because parents do not observe children's intelligence until after the child reaches at least 24 months of age. To justify this assumption, the authors cite psychology studies that conclude that parents cannot observe children's innate academic ability until the child is older than 24 months. This is a novel conceptual point, and is central to the paper's identification approach. The Philippines study also relies on the identifying assumption that the largest effects of child health on primary school outcomes in equation (1) are early in the first time period, that is from conception until the child is 24 months old, and on the assumption that changes in child health from 24 months of age until the start of primary school are not correlated with child health up to the age of 24 months. Finally, two implicit assumptions are that the impacts of parental education inputs before the child reaches primary school age (EI_1) and the impacts of current child health status (H_2) on primary school academic scores in equation (1) are negligible and thus can be dropped from that equation. These are relatively strong assumptions, but they deliver the needed econometric identification.

Together, these assumptions allow the authors to write the structural equation (1) as:¹¹

$$\begin{aligned}
 T_2 &= T(H_1, EI_2, \alpha, SC, YS) & (1') \\
 &= \beta_0 + \beta_1 H_1 + \beta_2 EI_2 + \beta_3 \alpha + \beta_4 SC + \beta_5 YS \\
 &= \beta_0 + \beta_1 \text{Height}_1 + \beta_2 EI_2 + \beta_3 \alpha + \beta_4 SC + \beta_5 YS
 \end{aligned}$$

where the second line is a simple linear approximation of the first line, and the third line explicitly uses child growth (measured by height) as the health indicator. That is, if good health leads to fast growth and poor health leads to slow growth, then H_1 is summarized by Height_1 (growth from conception until primary school). Equation (1') is for one child. Differencing across two siblings from the same family who attend the same school yields:

$$\Delta T_2 = \beta_1 \Delta \text{Height}_1 + \beta_2 \Delta EI_2 + \beta_3 \Delta \alpha + \beta_5 \Delta YS \quad (1'')$$

¹¹ The model in the paper has three time periods, from conception to 24 months (denoted in the paper as time period 0), from 24 months to 5 or 6 years (time period 1), and primary school age (time period 2), but the basic approach is the same as the description given here.

Equation (1'') is difficult to estimate because α and virtually all aspects of EI_2 are not observed and are likely to be correlated with the endogenous observed variables, ΔHeight_1 and ΔYS . One needs instrumental variables for ΔHeight_1 and ΔYS that are uncorrelated with $\Delta\alpha$ and ΔEI_2 , the differences in the innate intelligence and parental education inputs across the two siblings.

The authors use the differences in the dates of birth of the two siblings as the main instrument for ΔYS , which is arguably uncorrelated with $\Delta\alpha$ and ΔEI_2 (although the authors cannot completely rule out a story in which parents jointly plan birth spacing and the allocation of parental education inputs across siblings). Regarding ΔHeight_1 , the paper argues that the height of the older sibling by age 24 months is a valid instrument because it is uncorrelated with the α 's of both siblings (since neither is observed until after 24 months of age for the older sibling) and it has strong predictive power for ΔHeight_1 . Note that using instrumental variables also addresses the potential problem of bias due to measurement error in the height variables, if the measurement error is completely random.

Despite the innovative method of finding instruments for ΔHeight_1 and ΔYS in equation (1''), the estimation strategy remains open to several criticisms. The main problem with the estimation strategy is that it is not clear that the height of the older child at age 24 months, the instrument for ΔHeight_1 , is uncorrelated with differences in parental education inputs in the second time period (ΔEI_2). By the second time period of the older sibling, parents may take their children's health (which is measured by height) into account when making education input decisions. One could also quarrel with the implicit assumption that EI_1 does not have any direct effect on cognitive achievement. The existence of a multi-billion dollar industry in the U.S. claiming to boost infant intelligence (through "Baby Einstein" and related toys) suggests that many parents believe that EI_1 is valuable in improving later cognitive performance. This will introduce more variables into equation (1) for which instruments will be hard to find. The assumption that H_2 has no effect on child academic skills in period 2 (T_2) in equation (1) is also questionable. Overall, the approach used in the Philippines paper can be faulted, but the solutions to the criticisms raised here are far from obvious given the data at hand.

Using the estimation strategy explained above (modified to account for delayed enrollment and grade repetition), the Philippines study finds strong causal impacts of children's health status in the first two years of life (as measured by height at age 8) on several schooling outcomes. More specifically, better health leads to reductions in delayed enrollment, reduced

grade repetition, and greater learning per year of schooling as measured by test scores. The impacts appear to be large in that back of the envelope calculations based on the cost and impact (on child height) of an unrelated feeding program in India (Kielmann and Associates, 1983), together with the relationship between wages and education calculated from Philippines data, suggest that each dollar spent on a feeding program could provide a social return of at least three dollars and perhaps much more.

The third recent panel study that examines the impact of child health and nutrition on education is Alderman et al (2006). This paper estimates the impact of preschool height on years of completed schooling in a sample of 665 Zimbabwean young adults surveyed in early 2000. The authors also estimate impacts on delayed enrollment. The 2000 survey was a follow-up to two earlier surveys of children carried out in 1983/4 and 1987, and the authors appear to have had considerable tracking success: there is education data in the year 2000 for a remarkable 99% of the sample, including information obtained from relatives if the child had moved away from the study area. This is so in part because of the unusual nature of the sample: sample households all resided in resettlement communities, and these households had to renounce any land claims in other parts of Zimbabwe. Moreover, male adults were not allowed to out-migrate from the resettlement area. The high tracking rate is a noteworthy feature of these data and one that other studies should try to emulate.

The study uses a sibling comparison instrumental variables method related to that of Glewwe et al (2001). However, the source of variation that Alderman et al (2006) employ for their instrumental variables estimates is more exogenous than, and arguably an improvement over, previous (non-experimental) panel studies. Thus this quasi-experimental estimation approach is quite similar in spirit to the randomized evaluation studies described below in subsection 5.3.

Children in the sample were born between September 1978 and September 1986, an extremely volatile period both politically and in terms of living standards. In particular Zimbabwe in the late 1970s experienced the final years of a brutal civil war, and the country was later affected by back-to-back droughts in 1982-83 and 1983-84. The authors utilize variation in exposure to these large “shocks” across siblings while the children were 12 to 36 months of age to estimate the impact of preschool height on later outcomes, the underlying assumption being that children’s height (and their development more broadly) during that age range is more

sensitive to nutritional deficiencies than at other ages. The strong first stage relationships the authors estimate validate this view: exposure to these “shock” episodes during the key 12-36 month age range are strongly correlated with shorter child stature in earlier survey rounds (1983-84 and 1987). Over-identification tests confirm that these instruments do not blatantly violate the exclusion restriction. One concern is that schooling quality, or other unobserved inputs into education, are also affected by these macro-shocks, but there is no obvious reason why these shocks would translate into worse schooling quality several years later (when sample children enter primary school) only for the cohorts directly hit by the shock but not for their older or younger siblings.

Using this approach, Alderman et al. (2006) find that increased early childhood height is associated with significantly greater young adult height (in 2000) and more years of educational attainment. The effects are substantial: the increase from median child height in this sample to median height in the international reference (rich country) sample would lead to an additional 0.85 grades of completed schooling and over 3 cm in height. These fixed effects instrumental variable estimates are substantially larger than simple fixed effects estimates. The existence of data on schooling attainment, rather than just delayed school enrollment, makes it easier to translate these schooling effects into likely later impacts on income, using existing estimates of the returns to education. Note that estimated health effects for females and males are not significantly different from each other in the Zimbabwe study, a pattern also found in the experimental studies described below.

The estimates from Alderman et al (2006) thus largely confirm the findings of earlier cross-sectional and other panel studies in highlighting the important effect of early childhood height on later schooling outcomes. In summary, panel data provide additional possibilities for overcoming the estimation problems that plague studies based on cross-sectional data, but some estimation problems often remain. Undoubtedly, further data collection and innovative thinking will lead to improved estimates, like those in the Alderman et al (2006) study that utilize an arguably more convincing source of exogenous variation in child health, but the extent to which the remaining estimation problems can be resolved is difficult to predict.

On the other hand, it is worth noting that worries about estimation bias due to behavioral responses to health programs and policies may be exaggerated. Evidence in favor of this more optimistic viewpoint is found in a recent paper by Jacoby (2002) based on the same Philippines

data used by Glewwe et al. (2001). Jacoby found that parents did not reduce food given to their children at home in response to the availability of school feeding programs in Filipino primary schools. Even so, it would be imprudent to ignore the potentially serious estimation problems that arise in estimates based on nonexperimental cross-sectional and panel data. Thus the next subsection considers another approach: randomized evaluations.

5.3 Estimates based on randomized evaluations

Nutritionists and public health researchers have a long history of examining the impact of health programs and policies on cognitive and education outcomes using randomized evaluations. More recently, the difficulties of estimating the relationship between education outcomes and child health and nutrition have led some economists to initiate and analyze randomized evaluations in less developed countries. This subsection examines recent studies by both types of researchers, although it mainly focuses on the work of economists. Note also that this subsection does not review several recent studies that have used data from the Mexico PROGRESA project, mainly because that work is discussed in the chapter in this volume by Parker, Rubalcava, and Teruel (2007). An additional reason it does not cover this work is that the multiple components of PROGRESA assistance, and especially the income transfer component, in addition to health and nutrition interventions, complicate the task of isolating the impact of child health status per se on educational outcomes.

Many of the earliest randomized studies by nutritionists and other public health researchers focused on the impacts of specific nutrients that were lacking in children's diets. Studies in India and Indonesia by Soemantri, Pollitt, and Kim (1989), Soewwondo, Husaini, and Pollitt (1989), and Seshadri and Gopaldas (1989) found large and statistically significant impacts on cognitive development and school performance of iron supplementation among anemic children, but a study by Pollitt and others (1989) found no such impact in Thailand. See Nokes et al (1998) for a more complete survey of the iron supplementation literature.

Other studies have focused on parasitic infections, especially intestinal parasites. Kvalsig, Cooppan, and Connolly (1991) examined whipworms and other parasites in South Africa and found that drug treatments had some effect on cognitive and education outcomes, but some impacts were not statistically significant. Nokes and others (1992) evaluated treatment for whipworms in Jamaica and concluded that some cognitive functions improved from the drug

treatment, but others, particularly those related to academic performance in schools, appeared not to have changed substantially. Overall, the early experimental literature on the impact of treatment for intestinal parasites on child growth and cognition did not reach strong conclusions, as argued in the Dickson et al (2000) survey.

Other studies have focused on general food supplementation to supply calories and protein. The most well known of these is the INCAP study (Pollitt et al 1993, Martorell et al 1993) initiated in four Guatemalan villages in 1969, two of which were randomly selected to receive a porridge (*atole*) high in calories and protein while the other two villages received a drink (*fresco*) with less calories and no protein. Follow-up studies over the next two decades appear to show sizeable effects on later cognitive outcomes from providing the *atole* to mothers and young children.

These projects are arguably among the most convincing research to date showing long-term effects of childhood health and nutrition on later education, and on life outcomes more broadly.¹² Yet these studies are also subject to some criticisms. Many of these studies have relatively small sample sizes, such as 210 children in the South African study and 103 in the Jamaican study. Other studies (not reviewed here) include education interventions combined with health interventions, so the impact of the health intervention by itself cannot be credibly assessed.

The pioneering INCAP study is also open to some criticism. In one sense, it has a sample size of only four villages since the intervention did not vary within villages, and it is unclear if the existing studies fully account for the intracluster correlation of respondent outcomes in their statistical analyses, thus perhaps leading them to overstate the statistical significance of their findings. Second, strictly speaking, the control group also received an intervention, the *fresco* drink, albeit one with a relatively small benefit compared with what was received in the treatment group. Third, within each village receipt of the *atole* or *fresco* was voluntary, which implies that those who were treated were not a random sample of the population within each village. This means that the most convincing estimation strategy may be an intention to treat analysis, rather than direct estimation of the effect of child health on education. Finally, sample

¹² For a very recent summary of work done by nutritionists see the set of papers recently published in *The Lancet* (Grantham-McGregor et al., 2007; Walker et al., 2007; and Engle et al., 2007)

attrition is a major concern in the 1988-89 follow-up, as more than one quarter of the original sample were apparently lost, in sharp contrast to the exceptionally high tracking rate in the Alderman et al. (2006) study described above, or to very high tracking rates in other recent panel studies in less developed countries, most notably the Indonesia Family Life Survey (Thomas et al 2002). (Note that in recent work, the INCAP researchers have begun to extend their evaluation through 2002-03, see Behrman et al. 2003, but the analysis of long-run impacts remains preliminary at the time of writing this chapter.)

Three recent randomized evaluation studies by economists on the impact of health interventions on education outcomes are useful additions to this literature. These studies also evaluate actual interventions carried out by real-world non-governmental organizations (NGOs), and as such the findings of these studies may be of particular interest to policymakers in less developed countries. All three papers evaluate school-based health interventions which some have argued may be among the most cost-effective approaches for delivering health and nutrition services to children in less developed countries (Bundy and Guyatt 1996).

The first is that of Miguel and Kremer (2004), which evaluates a randomized program in Kenyan schools of mass treatment for intestinal worms using inexpensive deworming drugs. The study is based on a sample of 75 primary schools with a total enrollment of nearly 30,000 children, a much larger sample size than most other studies in this literature. The sampled schools were drawn from areas where there is a high prevalence of intestinal parasites among children. Worm infections – including hookworm, roundworm, whipworm and schistosomiasis – are among the most widespread diseases in less developed countries: recent studies estimate that 1.3 billion people worldwide are infected with roundworm, 1.3 billion with hookworm, 900 million with whipworm, and 200 million with schistosomiasis. Infection rates are particularly high in Sub-Saharan Africa (Bundy, et al. 1998; WHO 1993; see also the burden of disease figures in Table 2), where education outcomes and education progress are particularly low, as explained in Section 2. Geohelminths – hookworm, roundworm, and whipworm – are transmitted through poor sanitation and hygiene, while schistosomiasis is acquired by bathing in infected freshwater. School-aged children typically exhibit the greatest prevalence of infection and the highest infection intensity, as well as the highest disease burden, since morbidity is related to infection intensity (Bundy 1988). Recall from Table 2 that intestinal helminthes are estimated to account for about 3% of the total burden of disease among children age 5-14 in less

developed countries, and 4-5% in Sub-Saharan Africa. In fact, the impact of worms on the quality of life in the burden of disease calculations may be underestimated because they do not account for the impact of helminthes on education outcomes.

The educational impacts of deworming are considered a key issue in assessing whether the poorest countries should accord priority to deworming, but until recently research on these impacts has been inconclusive (see Dickson et al. 2000 for a survey). Indeed, earlier randomized evaluations on worms and education suffer from several important methodological shortcomings that may partially explain their weak results. Earlier studies randomized the provision of deworming treatment *within* schools to treatment and placebo groups, and then examine the impact of deworming on cognitive outcomes. However, the difference in educational outcomes between the treatment and placebo groups understates the actual impact of deworming if placebo group pupils also experience health gains due to local treatment externalities (due to breaking the disease transmission cycle). The earlier studies also failed to adequately address sample attrition, an important issue to the extent that deworming increases school enrollment.

The study by Miguel and Kremer finds that absenteeism in treatment schools was 25% (7 percentage points) lower than in comparison schools and that deworming increased schooling by 0.14 years per pupil treated (on average). This is a large effect given the low cost of deworming medicine; the study estimates an average cost of only US\$3.50 per additional year of school participation. The finding on absenteeism does not reflect increased school attendance on the part of children who attend school only to receive deworming drugs, since drugs were provided at only two pre-announced days per year, and attendance on those two days is not counted in the attendance analysis. There is no statistically significant difference in treatment effects between female and male students, echoing the finding discussed above in Alderman et al (2006).

Somewhat surprisingly, despite the reduction in absence no significant impacts were found on student performance on academic tests. It is unclear what exactly is causing this discrepancy, although one possibility is that the program led to more crowded classrooms and that this may have partially offset positive effects of deworming on learning in the treatment schools. In ongoing work, the authors of the Kenya study are collecting a new dataset, the Kenya Life Panel Survey (KLPS), in order to document the long-run impacts of the deworming program on educational attainment, cognitive skills, labor market outcomes, fertility, marital choices, health, physical strength and personal happiness.

The schooling data in Miguel and Kremer (2004) are noteworthy. School attendance was collected at sample schools by survey enumerators on unannounced days four to five times per year, rather than relying on school registers (which are thought to be unreliable) or on parent reports in household surveys, as done in most of the previous literature. Efforts were also made to follow children who transferred to other schools in the same Kenyan district. This yields a more detailed and reliable measure of school participation than the data available from most other studies. The Bobonis et al (2006) and Vermeersch and Kremer (2004) papers described below use similar measures of school attendance.

The authors found that child health and school participation – i.e., attendance, where dropouts are considered to have an attendance rate of zero – improved not only for treated students but also for untreated students at treatment schools (22% of pupils in treatment schools chose not to receive the deworming medicine) and for students at nearby primary schools located within 6 kilometers of treatment schools, with especially large impacts within 3 kilometers. The impacts on neighboring schools appear to be due to reduced disease transmission brought about by the intervention, an epidemiological externality. Econometric identification of the cross-school treatment spillovers on the worm infection rate relies on the randomized design of the project: conditional on the total local density of primary school pupils, there is random exogenous variation in the number of local pupils assigned to deworming treatment through the program. A key finding of the paper is that failure to take these externalities (or spillovers) into account would lead to substantial underestimation of the benefits of the intervention and the cost effectiveness of deworming treatment.

Bobonis, Miguel, and Sharma (2006) conducted a randomized evaluation in India of a health program that provided iron supplementation and deworming medicine to pre-school children age 2–6 years in 200 preschools in poor urban areas of Delhi. Even though only 30% of the sampled children were found to have worm infections, 69% of children had moderate to severe anemia according to international standards. After 5 months of treatment, the authors found large weight gains and a reduction of one-fifth in absenteeism, a treatment effect similar to the estimated school participation effect in the Miguel and Kremer (2004) study in Kenyan primary schools. The authors attempted to obtain estimates after 2 years, but high sample attrition and apparently non-random enrollment of new children into the preschools complicated attempts to obtain unbiased longer term impact estimates.

One plausible channel through which preschool attendance gains in Bobonis et al (2006) could have long-run impacts is an improvement in future primary school performance, and in fact, 71 percent of parents in the Indian study area claimed (in a baseline survey) that improved primary school preparedness was an important motivation for sending their own children to the preschools. There is some evidence linking preschool participation to later educational outcomes in both less developed and wealthy countries. Berlinski, Galiani, and Gertler (2006) find primary school test score improvements of 8 percent for children who had earlier participated in public preschool programs in Argentina. There is also evidence from the U.S. Head Start program that early childhood interventions reduce later grade repetition and increase educational attainment (Currie and Thomas 1995; Garces, Thomas, and Currie 2002). In terms of long-run evidence, Cascio (2004) finds 30 percent reductions in high school grade repetition among African-American and Latino children, and a 20 percent reduction among white children, who had earlier participated in public kindergarten programs in the U.S. South. Magnuson, Ruhm, and Waldfogel (2004) find evidence of medium-term gains from pre-kindergarten participation on first grade mathematics and reading, especially for children whose parents have low education or low income. Currie (2001) surveys the related U.S. literature and concludes that there is considerable evidence linking early childhood interventions to improvements in later educational attainment and cognitive development. It is possible that preschool attendance impacts could be even more persistent in less developed country contexts, where there are fewer school remedial programs and where households are poorer (consistent with the pattern in Magnuson et al 2004).

Another randomized evaluation using a similar research design is Vermeersch and Kremer (2004). Vermeersch and Kremer estimate the impact of a preschool feeding program in 50 Kenyan preschools. The daily feeding, with a protein enriched porridge, led to 30% higher preschool participation rates, and significant cognitive test score gains in schools with relatively experienced preschool teachers, although no significant cognitive gains in schools with less experienced teachers. The authors also document how the program led to large inflows of pupils into the feeding schools, suggesting that households' school choices may be sensitive to such programs. However, note that this feeding program is an order of magnitude more expensive than deworming treatment or micronutrient supplement, which will greatly increase the cost-benefit ratio.

Yet even these recent randomized evaluation studies have important limitations. The main puzzle with the Kenya deworming study is that increased school participation (primarily attendance, but also reduced dropping out) is not reflected in students' academic test scores or cognitive test scores. The authors present some cost-benefit analyses at the end of the paper that suggest that the intervention is cost-effective, but it is unclear exactly how to interpret these if the intervention does not increase learning of basic skills. Finally, since deworming treatment was found to affect child health in multiple ways – including lower intestinal worm load, reduced anemia, and (marginally) increased height-for-age – it is impossible to separately estimate the impact of each of these health improvements on education without imposing additional econometric structure.

The Bobonis et al. study (2006) encountered serious sample selection and attrition problems in the second year, which prevented a clear assessment of the long-term impact of the health intervention in India. It also does not present data on any type of child learning, and thus is limited to examining anthropometric outcomes and school enrollment and attendance. Finally, because all children received a combined treatment of iron supplements and deworming medicine, the India study cannot distinguish between the separate impacts of these two treatments. Vermeersch and Kremer (2004) are unable to distinguish between school attendance gains resulting from improved child nutrition per se versus a desire to receive food through the daily feeding program, which makes their estimates difficult to interpret relative to previous work (and a similar concern cannot be decisively ruled out in Bobonis et al (2006) with regard to the desire to receive more iron supplementation). A second limitation of Vermeersch and Kremer (2004) is the lack of anthropometric data on sample children, which limits comparability with previous studies in the literature.

6. Summary and concluding comments

This chapter has reviewed the most important estimation issues that complicate attempts to measure the impact of child health and nutrition status on education outcomes. As explained in Sections 3 and 4, the relationships between child health and schooling are very complex, and indeed there are multiple distinct relationships that are of potential interest, including the production function for academic skills, standard demand functions, and conditional demand functions. Perhaps the main message of this chapter is that it is very difficult, though not

impossible, to credibly estimate the relationship between child health and education. The two fundamental problems are the following: (1) it is impossible to obtain data on all variables that belong in the equations of interest, which raises serious problems of omitted variable bias; and (2) the variables that one does have data on are often measured with error, which can lead to problems of attenuation bias. These problems are not easy to fix, despite much richer data and the use of more careful estimation methods during the past ten years or so. Moreover, differences in data – in terms of both the health and education measures employed – complicates comparison of the magnitude of estimated health effects across studies.

Yet, despite these difficulties, most of the best recent studies using cross-sectional data, panel data, or data from randomized evaluations have found sizeable and statistically significant positive impacts of child health on education outcomes. Thus there is growing evidence of a causal impact of child health on education. There is no obvious reason to think that the litany of estimation problems described above systematically tend to overestimate the impacts of interest across all the different methodological approaches, data, and settings. A second noteworthy pattern emerging from the recent research is that there is no clear evidence of large gender differences in the impact of child health on education.

We close with a few suggestions for future research. In our view, future research on the links between child health and education outcomes should focus on two fronts (perhaps not surprisingly): better data, and better econometric identification. First, further analysis of panel data is warranted in both observational and randomized evaluation studies. Fortunately, more panel data collection efforts are now being undertaken in less developed countries than ever before, which will set the stage for such research. Improving sample tracking efforts will be critical to the success of ongoing studies, and the recent tracking success of the Alderman et al. (2004) and IFLS (Thomas et al 2002) studies means that sample attrition is not an insurmountable problem. Better panel datasets will also allow economists to directly estimate the long run impact of child health gains on their wages and living standards as adults, presumably the ultimate goal of much of this literature. Another area in which data can improve is in terms of the measurement of educational outcomes beyond simple parent reports on child school enrollment – perhaps following the approach in Miguel and Kremer 2004 – and of richer health outcomes. Health is multifaceted and it is unclear whether the height variables typically employed in this literature, while easy to measure, are really capturing the most critical

dimensions of health. Indeed, most of the diseases and health problems faced by children in developing countries, measured in terms of their contribution to the burden of disease as shown in Table 2, are unlikely to have strong impacts on height (the main exception being diarrhea), and variation in diseases and health problems across geographic regions is completely ignored by focusing on height.¹³

Second, more randomized evaluations should be conducted, especially by large international aid organizations. The results of these evaluations should be broadly disseminated, which will not be easy for these organizations because many studies will find that existing programs do not work as intended. Randomized studies should always compare their findings with standard cross-sectional or panel data estimates based on the control group data, making clear which of the three types of relationships discussed in section 3 are being compared. This will create a large source of information of the likely bias of non-experimental methods. It may be that there are many situations in which non-experimental methods do not suffer from substantial bias, but this will not become clear until a track record of results has been assembled. Randomized evaluations should also be designed in advance to go beyond the basic program impact evaluation results and by addressing broader theoretical and policy issues. Efforts to use structural modeling techniques in tandem with data from randomized evaluations (as in Todd and Wolpin 2003) are similarly a promising direction for this literature.

How large could the long-run effects of poor childhood health and nutrition on economic development really be? Unfortunately, the answer to this question remains elusive despite the recent research progress reviewed in this chapter. However, there is suggestive evidence from at least one once-developing country – the United States – that the long-run effects of public deworming investments could be very large indeed. Recent economic history research finds that the Rockefeller Sanitary Commission’s deworming campaigns in the United States South in the 1910s had major impacts on educational attainment and income (Bleakley 2002) and on agricultural productivity (Brinkley 1994). In fact, Bleakley (2002) estimates that each case of hookworm averted increased average school attendance by twenty percent. This historical evidence provides hope that current public health investments in children in less developed

¹³ Preliminary results from a recent randomized evaluation of the impact of providing eyeglasses to children with poor vision in a poor province in rural China (Glewwe, Park and Zhao, 2006) suggest large impacts on learning, but this is completely missed by focusing on child height.

countries could be planting the seeds for increased skills and, ultimately, greater prosperity during their adult lives.

References

- Alderman H, Behrman J, Lavy V, Menon R. Child health and school enrollment. *J Hum Resources* 2001; 36(1):185–205.
- Alderman H, Hoddinott J, Kinsey B. (2006). “Long Term Consequences of Early Childhood Malnutrition,” *Oxford Economic Papers*, 58(3), 450-474.
- Almond, Douglas, Kenneth Chay, and David S. Lee. (2006). The Cost of Low Birth weight, forthcoming *Quarterly Journal of Economics*.
- Angrist J, Imbens G, and Rubin D. Identification of Causal Effects Using Instrumental Variables, *Journal of the American Statistical Association* 1996; 91:444-472.
- Behrman J. The impact of health and nutrition on education. *World Bank Research Observer* 1996;11(1):25–37.
- Behrman J, Hoddinott J, Maluccio J, Quisumbing A, Martorell R, Stein, A. The Impact of Experimental Nutritional Interventions on Education into Adulthood in Rural Guatemala: Preliminary Longitudinal Analysis. 2003.
- Behrman, Jere R., and Mark R. Rosenzweig. (2004). Returns to Birth weight, *Review of Economics and Statistics*, 86(2), 586-601.
- Berlinski, Samuel, Sebastián Galiani, and Paul Gertler. 2006. “The Effect of Pre-Primary on Primary School Performance.” unpublished manuscript, Univ. of California, Berkeley.
- Bleakley, Hoyt. (2002). “Disease and Development: Evidence from Hookworm Eradication in the American South,” mimeo., MIT.
- Brinkley, Garland. (1994). “The Economic Impact of Disease in the American South, 1860-1940,” Ph.D. dissertation, University of California, Davis.
- Bobonis G, Miguel E, Sharma CP. (2006). “Iron deficiency, anemia and school participation,” *Journal of Human Resources*, 41(4), 692-721, .
- Bundy, D.A.P. (1988). “Population Ecology of Intestinal Helminth Infections in Human Communities.” *Philosophical Transactions of the Royal Society of London. Series B.* 321 (1207), 405-420.
- Bundy, D, Guyatt J. Schools for Health: Focus on Health, Education, and the School-Age Child. *Parasitology Today*; 12:1-16.
- Bundy, D.A.P., Chan, M-S., Medley, G.F., Jamison, D., and Savioli, L. (1998). “Intestinal Nematode Infections.” In *Health Priorities and Burden of Disease Analysis: Methods and*

Applications from Global, National and Sub-national Studies. Harvard University Press for the World Health Organization and the World Bank (in press).

- Cascio, Elizabeth U. 2004. "Schooling Attainment and the Introduction of Kindergartens into Public Schools." unpublished manuscript, University of California, Davis.
- Currie, Janet. 2001. "Early Childhood Education Programs." *Journal of Economic Perspectives* 15(2): 213-38.
- Currie, Janet, and Duncan Thomas. 1995. "Does Head Start Make a Difference?" *American Economic Review* 85(3): 341-64.
- Dickson R, Awasthi S, Williamson P, Demellweek C, Garner P. Effect of Treatment for Intestinal Helminth Infection on Growth and Cognitive Performance in Children: Systematic Review of Randomized Trials". *British Medical Journal*; 320, 1697-1701, 2000
- Engle, P., M. Black, J. Behrman, M. Cabral de Mello, P. Gertler, L. Capiiriri, R. Martorell and M. Young. 2007. "Strategies to Avoid the Loss of Developmental Potential in More than 200 Million Children in the Developing World". *Lancet* 369:229-242.
- Garces, Eliana, Duncan Thomas, and Janet Currie. 2002. "Longer-Term Effects of Head Start." *American Economic Review* 92(4): 999-1012.
- Glewwe, P. The Impact of Child Health and Nutrition on Education in Developing Countries: Theory, Econometric Issues, and Recent Empirical Evidence, *Food and Nutrition Bulletin*, 2005.
- Glewwe P, Jacoby H. Panel data. In: Grosh M, Glewwe P, eds. Designing household survey questionnaires for developing countries: lessons from 15 years of the living standards measurement study. New York: Oxford University Press, 2000.
- Glewwe P, Jacoby H. An economic analysis of delayed primary school enrollment in a low income country: the role of early childhood nutrition. *Rev Econ Stat* 1995;77(1):156-69.
- Glewwe P, Jacoby H, King E. Early childhood nutrition and academic achievement: a longitudinal analysis. *J Pub Econ* 2001;81(3):345-68.
- Glewwe, Paul, Albert Park and Meng Zhao. 2006. "The Impact of Eyeglasses on the Academic Performance of Primary School Students: Evidence from a Randomized Trial in Rural China." University of Minnesota and University of Michigan,
- Grantham-McGregor, Sally, Yin Bun Cheung, Santiago Cueto, Paul Glewwe, Linda Richter and Barbara Strupp. 2007. "Over Two Hundred Million Children Fail to Reach their Development Potential in the First Five Years in Developing Countries". *Lancet* 369:60-70.

- IAEEA (2000) “TIMSS 1999: International Mathematics Report,” International Study Center, Boston College.
- IAEEA (2003) “PIRLS 2001 International Mathematics Report,” International Study Center, Boston College.
- Jacoby H. Is there an intrahousehold flypaper effect? Evidence from a school feeding program. *Econ J* 2002;112(476):196–221.
- Kielmann, Arnfried A., and Associates. 1983. *Child and Maternal Health Services in Rural India: The Narangwal Experiment*, Baltimore, MD: The Johns Hopkins University Press.
- Kremer, Michael. (2003). “Randomized Evaluations of Educational Programs in Developing Countries: Some Lessons”, *American Economic Review: Papers and Proceedings*, 93 (2), 102-106.
- Kvalsig JD, Cooppan RM, Connolly KJ. The effects of parasite infections on cognitive processes in children. *Ann Trop Med Parasitol* 1991;73:501–6.
- Lee D. Trimming for Bounds on Treatment Effects with Missing Outcomes, Working Paper #51, Center for Labor Economics, University of California, Berkeley. 2002.
- Lopez, Alan, Colin Mathers, Majid Ezzati, Dean Jamison and Christopher Murray. 2006. “Global Burden of Disease and Risk Factors.” The World Bank. Washington, D.C.
- Magnuson, Katherine A., Christopher J. Ruhm, and Jane Waldfogel. 2004. “Does Prekindergarten Improve School Preparation and Performance?” *NBER Working Paper #10452*.
- Manski C. *Casual Inference in the Social Science*. Harvard University Press, Cambridge, 1995
- Martorell R, Habicht JP, Rivera JA. History and Design of the INCAP Longitudinal Study (1969-1977) and its Follow-up (1988-89). *Journal of Nutrition*; 125: 1027S-1041S, 1995.
- Miguel E, Kremer M. 2004. Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 2004;72(1):159–217.
- Nokes C, Grantham-McGregor S, Sawyer A, Cooper E, Bundy D. Parasitic helminth infection and cognitive function in school children. *Proceedings: Biological Sciences*, 247(1319):77–81. 1992.
- Nokes C, van den Bosch C, Bundy D. *The Effects of Iron Deficiency and Anemia on Mental and Motor Performance, Educational Achievement, and Behavior in Children: A Report of the International Nutritional Anemia Consultative Group*. USAID: Washington DC. 1998.

- Pollitt E, Gorman K, Engle P, Martorell R, Rivera J. Early supplemental feeding and cognition. Monographs of the Society for Research in Child Development, Serial No. 235, Chicago: University of Chicago Press. 1993.
- Pollitt E, Hathirat P, Kotchabhakadi N, Missel L, Valyasevi A. Iron deficiency and education achievement in Thailand. *Am J Clin Nutr* 1989(3);50:687–97.
- Seshadri S, Gopaldas T. Impact of iron supplementation on cognitive functions in preschool and school-aged children: the Indian experience. *Am J Clin Nutr* 1989;50(3):675–86.
- Soemantri AG, Pollitt E, Kim I. Iron deficiency anemia and education achievement. *Am J Clin Nutr* 1989; 50(3):698–702.
- Soewondo S, Husaini M, Pollitt E. Effects of iron deficiency on attention and learning processes of preschool children: Bandung, Indonesia. *Am J Clin Nutr* 1989;50(3):667–74.
- Thomas D, Frankenberg E, Smith J. Lost but Not Forgotten: Attrition and Follow-up in the Indonesia Family Life Survey. *Journal of Human Resources*; 36: 556-592, 2002.
- Todd P, Wolpin K. Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility: Assessing the Impact of a School Subsidy Program in Mexico”, PIER Working Paper #03-022, 2003.
- UNESCO (2002) *Education for All: Is the World On Track?* UNESCO Publishing, Paris.
- UNESCO (2003) “UNESCO Institute of Statistics Electronic Database”
http://www.uis.unesco.org/ev.php?URL_ID=5187&URL_DO=DO_TOPIC&URL_SECTION=201
- United Nations ACC/SCN (in collaboration with the International Food Policy Research Institute). 4th Report on the World Nutrition Situation. Washington DC, 2000.
- Vermeersch C, Kremer M. School Meals, Educational Achievement and School Competition: Evidence from a Randomized Evaluation. Unpublished working paper, World Bank and Harvard University.
- Walker, S. T. Wachs, J. Meeks Gardener. B. Lozoff, G. Wasserman, E. Pollitt and J. Carter.. 2007. “Child Development: Risk Factors for Adverse Outcomes in Developing Countries”. *Lancet* 369:145-157.
- World Bank. World development report 2000/2001: attacking poverty. New York: Oxford University Press, 2001.
- World Bank. 2005 World development Indicators. Washington, DC, 2005.
- World Health Organization. (1993). The Control of Schistosomiasis. Second Report of the WHO

Expert Committee. (Technical Report Series 830). WHO, Geneva.

World Health Organization. The World Health Report 2000. Geneva, 2000.

Table 1: Child Health and Nutrition in Developing Countries

<u>Region</u>	<u>Underweight (%)</u> <u>(children < 5 years)</u>	<u>Under-Five Mortality</u> <u>Rate (per 1000)</u>	<u>Under-Five Mortality</u> <u>Rate (per 1000)</u>
	<u>1995-2003</u>	<u>1990</u>	<u>2003</u>
East Asia and Pacific	15	59	41
Latin America	9	53	33
Middle East/N. Africa	15	77	53
South Asia	48	130	92
Sub-Saharan Africa	-	187	171
All Developing Countries	27	103	87

Source: World Bank (2005)

Table 2: Estimated Burden of Disease for Children in Developing Countries, 2001

Children Age 0-4 Years

	<u>All Less</u> <u>Developed</u> <u>Countries</u>	<u>East</u> <u>Asia and</u> <u>Pacific</u>	<u>Latin</u> <u>America</u> <u>and</u> <u>Caribbean</u>	<u>Middle</u> <u>East and</u> <u>North</u> <u>Africa</u>	<u>South</u> <u>Asia</u>	<u>Sub-</u> <u>Saharan</u> <u>Africa</u>
Communicable Diseases, of which:	52.4%	30.8%	23.6%	35.1%	48.7%	72.6%
AIDS	2.5%	0.3%	0.8%	0.1%	0.3%	6.0%
Diarrhea	12.6%	10.8%	8.0%	11.1%	14.5%	12.8%
Pertussis	2.7%	0.9%	1.5%	1.5%	2.7%	3.8%
Measles	4.0%	2.1%	0.0%	1.4%	3.1%	6.6%
Tetanus	1.3%	0.7%	0.1%	0.5%	1.8%	1.6%
Malaria	8.5%	1.4%	0.3%	2.7%	1.3%	20.2%
Respiratory Infections	14.8%	10.5%	7.5%	7.5%	18.7%	14.7%
Perinatal	21.1%	28.8%	26.7%	19.8%	26.3%	12.4%
Nutrition Problems	4.4%	4.5%	3.7%	5.8%	4.3%	4.1%
Noncommunicable Illnesses:	18.6%	30.3%	41.8%	33.4%	17.3%	8.4%
Injuries	3.6%	5.5%	4.1%	5.9%	3.3%	2.5%
Percent of Healthy Years Lost	15.1%	8.4%	8.6%	11.3%	16.8%	28.6%

Children Age 5-14 Years

Communicable Diseases, of which:	36.9%	23.0%	16.6%	15.0%	38.2%	56.4%
Tuberculosis	1.5%	1.4%	0.7%	0.5%	1.7%	1.9%
AIDS	3.7%	0.1%	1.4%	0.0%	0.4%	11.3%
Diarrhea	1.6%	2.1%	3.7%	1.7%	1.2%	1.2%
Measles	6.5%	4.7%	0.0%	2.9%	6.5%	10.4%
Tetanus	1.3%	0.8%	0.0%	0.2%	2.0%	1.4%
Malaria	1.2%	0.5%	0.3%	0.8%	0.9%	2.3%
Intestinal helminths	2.7%	3.3%	1.7%	1.7%	1.5%	4.4%
Respiratory Infections	6.6%	5.6%	4.2%	4.1%	6.0%	9.1%
Nutrition Problems	4.1%	4.9%	2.1%	7.1%	4.4%	2.8%
Noncommunicable Illnesses:	30.3%	38.4%	58.4%	40.9%	29.1%	15.5%
Injuries	28.5%	33.6%	22.6%	36.7%	28.1%	25.0%
Percent of Healthy Years Lost	0.8%	0.5%	0.5%	0.6%	0.9%	1.4%

Source: World Bank (2006, Table 3C).

**Table 3. Primary School Gross Enrollment Rates
(percent of students of primary school age)**

<u>Area</u>	<u>1960</u>	<u>1970</u>	<u>1980</u>	<u>1990</u>	<u>2000</u>
<u>Country group</u>					
Low-income	65	77	94	102	102
Middle-income	83	103	101	103	110
High-income	109	100	101	102	102
<u>Region</u>					
Sub-Saharan Africa	40	51	80	74	77
Middle East/North Africa	59	79	89	96	97
Latin America	91	107	105	106	127
South Asia	41	71	77	90	98
East Asia	87	90	111	120	111
East Europe/Former Soviet Union (FSU)	103	104	100	98	100
Organization for Economic Cooperation and Development (OECD)	109	100	102	103	102

Note: Countries with populations of less than 1 million are excluded.

Sources: Barro and Lee data set; UNESCO (2002); World Bank (2003)

Table 4. Primary School Enrollment, Repetition, and Grade 4 Survival Rates (percents)

<u>Areas</u>	<u>Gross</u> <u>enrollment</u> <u>2000</u>	<u>Net</u> <u>enrollment</u> <u>2000</u>	<u>Repetition</u> <u>2000</u>	<u>On-time</u> <u>enrollment</u> <u>2000</u>
<u>Country group</u>				
Low-income	102	85	4	55
Middle-income	110	88	10	61
High-income	102	95	2 ^a	73 ^b
<u>Region</u>				
Sub-Saharan Africa	77	56	13	30
Middle East/North Africa	97	84	8	64
Latin America	127	97	12	74
South Asia	98	83	5	-
East Asia	111	93	2	56
East Europe/FSU	100	88	1	67 ^a
OECD	102	97	2 ^a	91 ^a

Notes: Countries with populations of less than 1 million are excluded.

a. Data are based on between 25-50% of the total population of the country group or region.

b. Data are based on between 10-25% of the total population of the country group or region.

Source: UNESCO (2003)

Table 5. Secondary School Gross Enrollment Rates
(percent of students of secondary school age)

<u>Area</u>	<u>1960</u>	<u>1970</u>	<u>1980</u>	<u>1990</u>	<u>2000</u>
<u>Country group</u>					
Low-income	14	21	34	41	54
Middle-income	21	33	51	59	77
High-income	63	74	87	92	101
<u>Region</u>					
Sub-Saharan Africa	5	6	15	23	27
Middle East/North Africa	13	25	42	56	66
Latin America	14	28	42	49	86
South Asia	10	23	27	39	47
East Asia	20	24	44	48	67
East Europe/FSU	55	64	93	90	88
OECD	65	77	87	95	107

Notes: Countries with populations of less than 1 million are excluded.

Source: Barro and Lee data set; UNESCO (2003); World Bank (2003)

Table 6. Gender Disparities in Gross Primary and Secondary Enrollment Rates, 

<u>Area</u>	<u>Primary</u>		<u>Secondary</u>	
	<u>Boys</u>	<u>Girls</u>	<u>Boys</u>	<u>Girls</u>
<u>Country group</u>				
Low-income	107	98	60	47
Middle-income	112	108	77	78
High-income	102	101	100	102
<u>Region</u>				
Sub-Saharan Africa	83	71	29	24
Middle East/North Africa	101	92	71	61
Latin America	129	125	83	89
South Asia	107	90	53	39
East Asia	112	111	73	60
East Europe/FSU	100	99	88	89
OECD	102	102	106	108

Notes: Countries with populations of less than 1 million are excluded.

Source: World Bank (2003)

Table 7. Mean Mathematics and Reading Achievement, TIMSS and PIRLS Studies

<u>Country</u>	<u>Mathematics (TIMSS) 1999</u>		<u>Reading (PIRLS) 2001</u>
	<u>Grade 7</u>	<u>Grade 8</u>	<u>Grade 4</u>
France	-	-	525
Japan	-	579	-
U.K. (England)	-	-	553
U.S.	-	502	542
Argentina	-	-	420
Belize	-	-	327
Chile	-	392	-
Colombia	-	-	422
Indonesia	-	403	-
Iran	-	422	414
Jordan	-	428	-
Korea (South)	-	587	-
Kuwait	-	-	396
Malaysia	-	519	-
Morocco	337	-	350
Philippines	345	-	-
South Africa	-	275	-
Thailand	-	467	-
Tunisia	-	448	-
Turkey	-	429	449

Source: IAEEA (2000, 2003)