Chapter 55

SCHOOLING IN DEVELOPING COUNTRIES: THE ROLES OF SUPPLY, DEMAND AND GOVERNMENT POLICY*

PETER F. ORAZEM

Department of Economics, Iowa State University, 267 Heady Hall, Ames, IA 50011-1070, USA

ELIZABETH M. KING

The World Bank, 1818 H Street, NW Washington, DC 20433, USA

Contents

Abstract 3476
Keywords 3477
1. Introduction 3478
2. Costs, returns and schooling gaps 3479
  2.1. A model of schooling length 3481
  2.2. Male–female schooling gaps 3483
  2.3. Urban–rural schooling gaps 3492
3. How do government policies affect schooling gaps? 3496
  3.1. A local schooling market 3497
  3.2. Educational investments without government intervention 3498
  3.3. Modeling different interventions 3499
    3.3.1. Subsidies of schooling costs 3500
    3.3.2. Vouchers 3500
    3.3.3. Unconditional income transfers 3502
    3.3.4. Conditional income transfers 3502
  3.4. Endogenous program placement and participation 3504
    3.4.1. Application: Evaluating voucher programs in Latin America 3504
4. Static model 3507
  4.1. A household model 3507

* This document represents the opinions of the authors and does not represent the opinions of the World Bank or its Board of Directors. Orazem began work on this paper while serving as Koch Visiting Professor of Business Economics at the University of Kansas, School of Business. We are grateful to Suzanne Duryea at the Inter-American Development Bank for early discussions about measuring education returns, and to Caridad Araujo, Arjun Bedi, Victoria Gunnarsson, Claudio Montenegro, and Norbert Schady for providing access to unpublished data used in the review. Esther Duflo and T. Paul Schultz provided extensive comments on earlier drafts that greatly improved the paper.

Handbook of Development Economics, Volume 4

DOI: 10.1016/S1573-4471(07)04055-7

4.2. Estimating the static model 3511
4.2.1. Application: The role of missing ability in evaluations of nutrition on cognitive attainment 3513
4.3. Estimating schooling demand 3515
4.3.1. Application: Using variation in school supply to identify years of schooling completed and returns to schooling 3515
4.4. Discussion 3520
5. Dynamic models 3521
5.1. Application: Structural estimation of the dynamic model: Mexico's PROGRESA Program 3524
5.2. Dynamic labor supply equations 3526
5.2.1. Application: Jacoby and Skoufias: Enrollment responses to anticipated and unanticipated income shocks 3527
5.2.2. Intergenerational wealth transfers 3529
5.2.3. Unconstrained households 3530
5.2.4. Constrained households 3530
5.2.5. Implications for studies of intergenerational transfers for poor and wealthy households 3531
5.2.6. Genetic transfers of human capital 3532
5.2.7. Application: The 1998 Indonesia Currency Crisis 3533
6. Measurement matters 3535
6.1. Measurement of time in school 3535
6.1.1. Enrollment status 3536
6.1.2. Years of schooling 3537
6.1.3. Grade-for-age 3538
6.1.4. Attendance: Days or hours of school 3540
6.2. Measurement of school outcomes 3541
6.2.1. School quality 3545
6.2.2. Promotion, repetition and dropping out 3545
6.3. Measurement of exogenous variables 3546
6.3.1. Direct price of schooling 3546
6.3.2. Opportunity cost of child time 3548
7. Conclusions 3549
References 3553

Abstract

In developing countries, rising incomes, increased demand for more skilled labor, and government investments of considerable resources on building and equipping schools and paying teachers have contributed to global convergence in enrollment rates and completed years of schooling. Nevertheless, in many countries substantial education gaps persist between rich and poor, between rural and urban households and between males and females. To address these gaps, some governments have introduced school vouchers or cash transfers programs that are targeted to disadvantaged children. Others have initiated programs to attract or retain students by expanding school access or

by setting higher teacher eligibility requirements or increasing the number of textbooks per student. While enrollments have increased, there has not been a commensurate improvement in knowledge and skills of students. Establishing the impact of these policies and programs requires an understanding of the incentives and constraints faced by all parties involved, the school providers, the parents and the children.

The chapter reviews the economic literature on the determinants of schooling outcomes and schooling gaps with a focus on static and dynamic household responses to specific policy initiatives, perceived economic returns and other incentives. It discusses measurement and estimation issues involved with empirically testing these models and reviews findings.

Governments have increasingly adopted the practice of experimentation and evaluation before taking steps to expand new policies. Often pilot programs are initiated in settings that are atypically appropriate for the program, so that the results overstate the likely impact of expanding the program to other settings. Program expansion can also result in general equilibrium feedback effects that do not apply to isolated pilots. These behavioral models provide a useful context within which to frame the likely outcomes of such expansion.

Keywords

education, household demand for education, education policy

JEL classification: I20, I21, I28, D13

1. Introduction

Enrollment rates and years of schooling have risen in most countries, a result of successive generations of parents investing in children's education. Over time, these investments have narrowed the differences in schooling across and within cohorts of children, across and within countries, and between and within genders. In 1960, the average schooling of men aged 25 and over in advanced countries was 5.8 times that of men in developing countries; in 2000, this ratio was down to 2.4.[1] During the same period, in developing countries, women's average schooling level as a ratio of men's increased from 0.5 to 0.7. While increasing incomes, shifts in demand for more skilled labor, and more classrooms have contributed to some global convergence in education as measured by years of schooling, substantial education gaps persist, however, such as between rural and urban households and also between males and females, in some settings. These gaps lead to the questions, what are the sources of these gaps and can they be influenced by economic growth, government policy, or international pressure?

Governments devote widely different shares of their budgets to education, at a range of 6–25 percent in 2000 across African countries alone. Parents also devote considerable resources to investments in their children, but also with high variability – in 2001, from 6 percent of total (public and private) spending for primary and secondary education in India to 33 percent in the Philippines. For the poorest parents who send their children to school, such investments will be most, if not all, of the wealth they transfer to their children.

Despite this large variation in rates of human capital investment, estimated private rates of return to years of completed schooling are remarkably similar across countries and across sub-populations within countries. The estimated proportional increase in labor earnings per year of schooling across many developing countries averaged 8% with an interquartile range of 5–10%. In comparison, there is less agreement about the magnitude of social rates of return to schooling, perhaps due to less agreement about how to measure these returns. Nonetheless, there seems broad agreement that schooling benefits society in many ways – in terms of better infant, child and maternal health; reduced fertility; enhanced ability to adopt new technologies or to cope with economic shocks; and rising labor productivity and sustainable economic growth. Reflecting the large collective evidence on these returns, theoretical models have included investments in human capital as an important source of persistent economic growth.

Governments have invested considerable resources on education. Numerous initiatives have been attempted aimed at increasing the returns to those investments. Some initiatives have aimed at raising school quality such as setting higher eligibility requirements for teachers or increasing the number of textbooks in the hands of students. Remedial programs have tried to reduce dropout rates. More recently, initiatives have

[1] Data as given in http://www.worldbank.org/edstats; raw data for these years are based on UNESCO statistics.

attempted to increase attendance at existing schools through school vouchers or cash transfers conditioned on child enrollment. Relatively few studies have found empirical evidence compelling enough to merit continued support for these initiatives. Nevertheless, governments often implement policy initiatives on a national scale straight from concept. Alternatively, initiatives are introduced primarily in settings where they are expected to be atypically successful rather than testing how they might perform in more diverse and challenging areas. The practice of experimentation and evaluation before policy adoption is a welcome recent innovation in many settings.

This chapter examines the magnitude of schooling gaps between population groups, why gaps occur, why they persist or diminish over time, how they respond to economic shocks, and how they are transmitted from parent to child. Models can assist us in forecasting where policies are most likely to succeed and where they are likely to fail. Behavioral models help to explain why experimental outcomes may be favorable in some settings and unfavorable in others, and for identifying the most promising locations where experiments can be replicated. Finally, behavioral models help to structure the empirical measurement of responses to government policies and economic circumstances that guide our ability to forecast how households react to government human capital investments and to perceived economic returns.

The chapter opens with a review of patterns, trends, and explanations of urban, rural, male and female education levels in the developing world (Section 2). Next, the chapter presents a model of how government policies affect education levels (Section 3). Static and dynamic models used to guide empirical studies of educational choices follow (Sections 4–5). Estimating these models involve a host of measurement issues which are addressed in Section 6.

2. Costs, returns and schooling gaps

We have already mentioned that education trends in the past half-century have been characterized paradoxically by greater convergence and also by persistent gaps. With or without economic growth, the opening up of national borders, faster information and communication technology, and international social mandates coupled with aid have raised the demand for education, even in poor countries. Indeed, many poor countries do foster high enrollments, especially in their urban areas – but other poor countries do not. Within countries, some groups are faster than others to respond to improved school access, thus widening within-country variation in schooling. Before turning to the factors that produce or sustain education gaps even as global factors appear to support convergence, we illustrate the different across- and within-country patterns in education levels of developing countries using age-enrollment "pyramids" for two low-income countries (Ethiopia and Tanzania) and two lower middle-income countries (Morocco and Turkey). The pyramids show markedly different patterns in the proportion of children enrolled in school by age, sex, and urban or rural residence (Fig. 1).

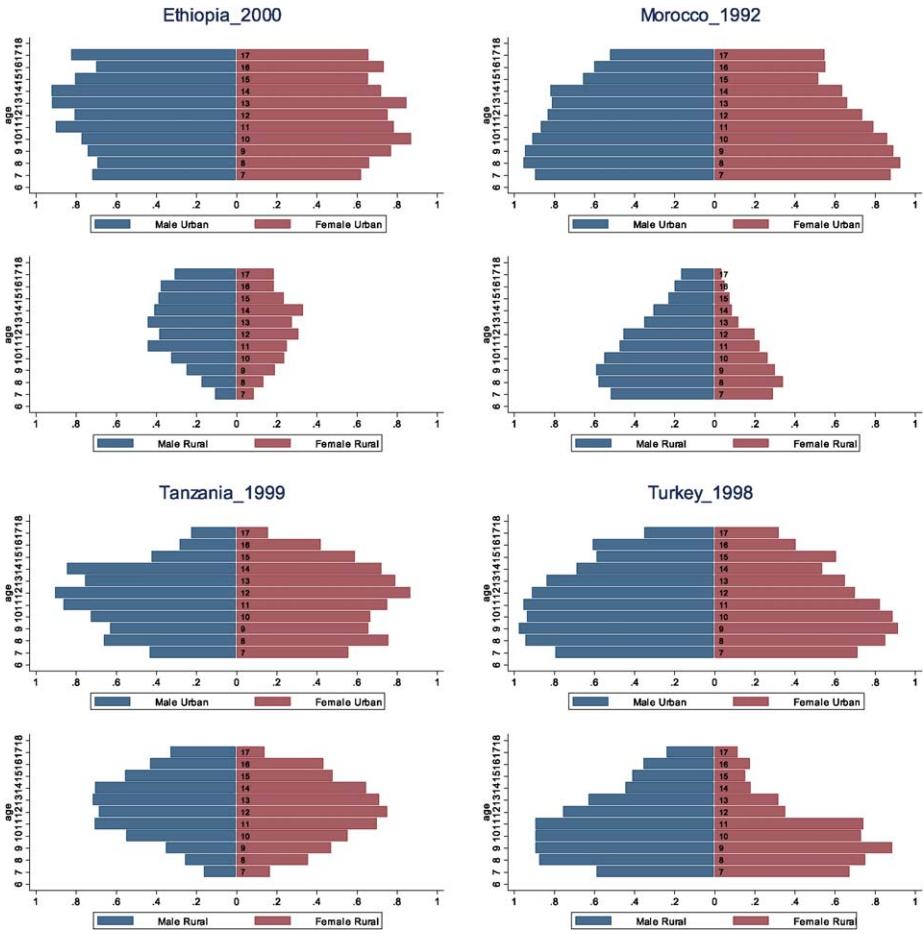

Figure 1. Age-enrollment bars, by gender and residence, selected countries.

Ethiopia is the poorest of the countries but has a remarkable proportion of its urban boys and girls enrolled in school and, through age 13, urban boys and girls attend school in equal proportions. In rural areas, however, enrollment rates never exceed 50 percent for any age or gender. Most urban children enter school at age 7, but rural children typically delay entry if they enter at all, leading to a large education gap between urban and rural children. Unlike urban areas, rural girls receive significantly less schooling than do rural boys.

Tanzania is only modestly wealthier than Ethiopia but has very different enrollment patterns. The gap between urban and rural schooling is much smaller, partly because rural children are much more likely to attend school and partly because urban children attend less. At the oldest ages, the rural children are more likely to be in school than in

Ethiopia, although this could be reflecting higher rates of grade repetition rather than more grades attained. There are no substantial schooling gaps between boys and girls except at these older ages.

Despite being much wealthier, Morocco looks more like Ethiopia than does Tanzania. The Moroccan age-enrollment pyramid is widest at its base, reflecting the more typical pattern of early entry into school and increasing dropout rates at older ages. Urban boys and girls receive similar schooling, but rural girls receive much less schooling than do boys. Turkey, the wealthiest of the four countries, share with the other three countries the pattern of enrollment rates dropping sharply once children reach 13 in both urban and rural areas, with a particularly pronounced dropout rate in rural areas. Girls and boys are treated similarly in urban areas, but rural girls drop out more rapidly after age 11.

2.1. A model of schooling length

Underlying the gaps within each country are individual household decisions of how long to send their children to school. Those decisions reflect trade-offs between schooling costs in the present against anticipated larger earnings capacity in the future. Building on Rosen's (1977) formulation, we model individual earnings as a function of the individual's stock of human capital upon leaving school, $q = q(E, z)$. Human capital production depends positively on E: years of schooling; and positively on z: a vector of exogenous factors that also raise school productivity such as ability or school quality so that $q_{Ez} > 0$. Ignoring the direct costs of schooling, the gross return per year of schooling is defined as $\rho = \frac{\mathrm{d}q}{\mathrm{d}E}\frac{1}{q} = \frac{q_E}{q}$. Schooling is subject to diminishing returns, and so ρ is assumed to diminish in E.

Schooling is not without cost, however. We assume that p is a constant cost per unit of schooling, and that this cost includes exogenous schooling tuition and other fees plus a fixed opportunity cost of time. Because costs are incurred and returns are earned over a period of time, we discount back to the initial period using an exogenous interest rate $r(y)$, $r'(y) < 0$, where y is a measure of household income. We assume that wealthier families have better access to credit markets and can command better credit terms; thus, interest rates are presumed to decrease in y.

The lifetime discounted value of income at birth net of schooling costs will be

$$\begin{aligned} V_0(E) &= \int_E^N q(E, z)\mathrm{e}^{-rt}\,\mathrm{d}t - \int_0^E p\mathrm{e}^{-rt}\,\mathrm{d}t \\ &= \frac{q}{r}\left(\mathrm{e}^{-rE} - \mathrm{e}^{-rN}\right) + \frac{p}{r}\left(\mathrm{e}^{-rE} - 1\right) \end{aligned} \tag{2.1}$$

where N is the anticipated retirement age. The formulation presumes that the student devotes full time to school for a period of length E and thereafter devotes full time to the labor market earning a wage $q(E, z)$ for a work career spanning from period E to N. At birth, N is very large, and so e^{-rN} approaches zero. Consequently, lifetime discounted

income can be approximated by

$$V_0(E) = \frac{q}{r}\mathrm{e}^{-rE} + \frac{p}{r}\left(\mathrm{e}^{-rE} - 1\right). \tag{2.2}$$

The optimal length of time to spend in school is selected so as to maximize discounted lifetime income at birth. Taking the derivative of $V_0(E)$ with respect to E and setting the result equal to zero, we obtain $V_0'(E) = -q\mathrm{e}^{-rE} + \frac{q_E}{r}\mathrm{e}^{-rE} - p\mathrm{e}^{-rE} = 0$. Rearranging, this reduces to

$$r + r\frac{p}{q} = \frac{q_E}{q} = \rho. \tag{2.3}$$

The student will remain in school until the gross rate of return is equated with the interest rate plus a term that rises in the cost of schooling. If interest rates are higher in developing countries, we would expect returns to schooling to be higher and length of time in school to be lower in the poorest countries, other things constant.

The relationship (2.3) implicitly defines years of schooling as a function of the exogenous variables

$$E = E(y, p, z). \tag{2.4}$$

Years of schooling rise in household income (which lowers the interest rate) and fall as schooling becomes more costly, while the impact of z on years of schooling is uncertain.[2] The gross rate of return to schooling will also be endogenous, determined by the same factors so that $\rho = \rho(y, p, z)$.

In the typical Mincerian (1974) formulation, direct costs of schooling are assumed to be zero. In that case, the first-order condition simplifies to $r = \frac{q_E}{q} = \rho$. Inserting $q = \frac{q_E}{r}$ back into $V_0(E)$, and imposing $p = 0$ and $\mathrm{e}^{-rN} = 0$, we get the relationship $rV_0(E) = (\frac{q_E}{r})\mathrm{e}^{-rE}$. Rearranging, this yields a log linear relationship between earnings and years of schooling, $\ln q = \ln(rV_0(E)) + rE$. The first-order condition implies that the coefficient of years of schooling E is equal to the gross rate of return ρ when $p = 0$. When costs are not zero, the relationship is the more unwieldy, $\ln(q + \frac{p}{r}) = \ln(rV_0(E) + \frac{p}{r}) + rE$. When $p > 0$, the first-order condition implies that the coefficient of years of schooling E will be $r < \rho$.

Thus far, we have equated human capital with the number of years of schooling that students spend in school. As we will discuss later in the chapter, due to grade repetition, student absences and variation in school quality, number of years of schooling is not a perfect measure of human capital accumulation. Other studies use an even simpler measure, enrollment in a particular grade or at a certain age. In settings where nearly a half of the population has never been to school or reports zero years of schooling, such as the case of rural Ethiopia, the most important schooling decision may indeed be whether a

[2] To derive the comparative static effects, it is convenient to rewrite (2.3) as $-q + \frac{q_E}{r} - p = 0$. Letting $A = (\frac{1}{r}q_{EE} - q_E) < 0$; $\frac{\mathrm{d}E}{\mathrm{d}p} = (1/A) < 0$; $\frac{\mathrm{d}E}{\mathrm{d}y} = r'(y)(\frac{q_E}{Ar^2}) > 0$; and $\frac{\mathrm{d}E}{\mathrm{d}z} = (q_z - \frac{q_{Ez}}{r})/A$ which has an ambiguous sign.

child ever attends school or not and so enrollment is the appropriate measure. In contexts where nearly all children of school age enter school, enrollment is an incomplete measure of the household's schooling decisions unless the focus of the study is on the timing of school entry. Yet other studies consider the quality, not just the quantity, of schooling as endogenous rather than exogenous, with households or individuals choosing which schools to attend and students deciding on their level of effort. Indeed, as data on measures of academic performance have become more available, more studies have turned to learning as measured by test scores as a better indicator of human capital.[3] We return to these measurement issues in Section 6.

The relationship in (2.3) suggests that two students of similar abilities facing identical interest rates, schooling costs and human capital production processes will choose the same length of time in school. However, gaps in schooling attainment exist between men and women in many developing countries, most often favoring men. Even larger gaps exist between urban and rural populations, nearly always favoring urban residents.

2.2. Male–female schooling gaps

Current enrollment rates for children and years of schooling completed for adults show gender gaps, but overall, women in developing countries have gained relative to men with respect to education.[4] These patterns emerge from information from the most recently available household surveys (e.g., country censuses, Living Standards Measurement Surveys, and Demographic and Health Surveys) for 70 developing countries; the data are weighted to produce country-wide averages. Figure 2(a) and 2(b) illustrates the range of gender enrollment gaps for two age groups 7–11 and 15–17 in urban and rural areas; the 12–14 year-old group had plot patterns lying between these two. By differentiating between urban and rural areas at the same time, we see an important aspect of the pattern in gender gaps.

Countries plotted in the northeast quadrant have enrollment gaps favoring boys in both urban and rural areas, while those in the southwest quadrant have gaps favoring girls. Countries in the southeast quadrant have gaps favoring girls in rural areas and boys in urban areas, whereas those in the northwest quadrant have gaps favoring girls in urban areas and boys in rural areas. Countries above the 45° line have more positive male–female gaps in urban areas while those below the 45° line have more positive male–female gaps in rural areas. A box centered on (0, 0) with sides of length 0.2 helps to illustrate which gaps are larger than 10% in either direction. Any point lying outside the box indicates at least one gap larger than 10%.

[3] In turn, several studies have estimated the return to the quality of schooling as separate from the return to years of schooling (e.g., Moffitt, 1996; Altonji and Dunn, 1996; Case and Yogo, 1999; Bedi and Edwards, 2002).

[4] The average years of schooling attained is defined as highest grade completed rather than the actual number of years enrolled in school. Due to grade repetition, the highest grade attained can imply fewer years of schooling than the number of years actually spent in school. We have no separate information on grade repetition from the surveys.

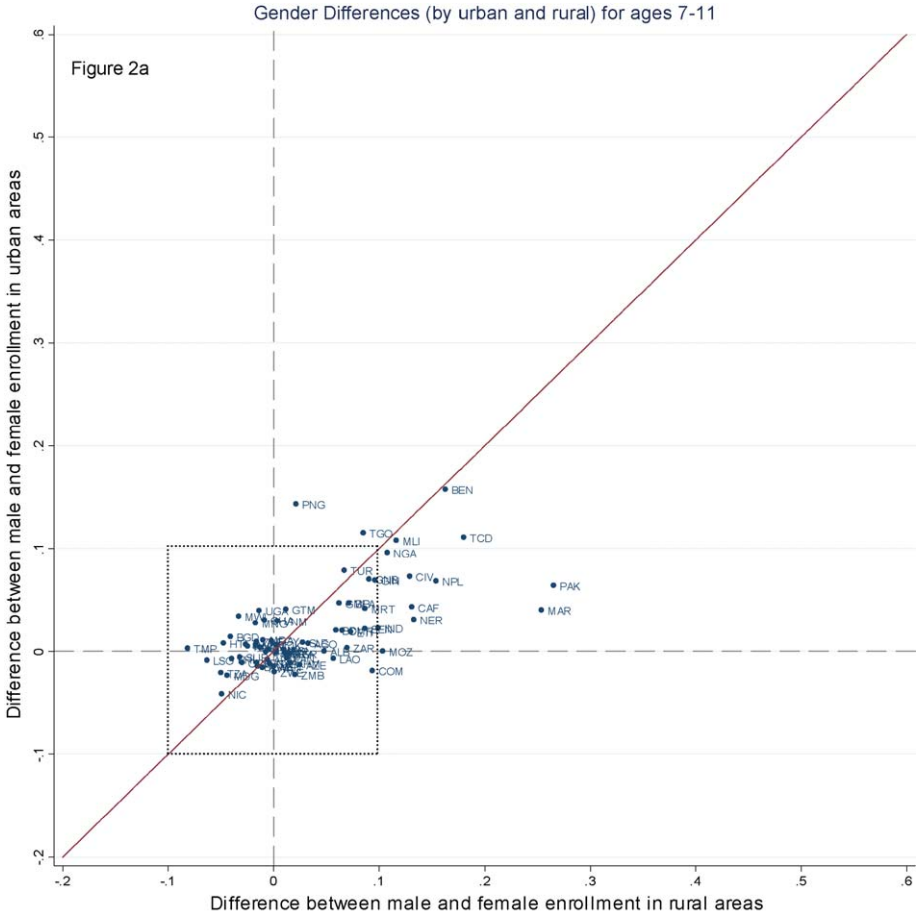

Source: Graphs computed from data from latest household surveys in 70 countries; database prepared for the World Bank, *World Development Report 2007*. All data, with few exceptions, are from year 2000 or later.

Figure 2. Combinations of male–female differences in enrollment rates, by urban and rural residence, age and country.

Several stylized facts emerge from the plots in panels 2a and 2b.

(1) In the youngest age group (7–11), male–female gaps tend to be small in both urban and rural areas. The gaps favor girls in many countries, but those differences tend to be small. The largest of gaps favor boys, most in rural areas.

(2) As children age, the variance in gaps increases. By ages 15–17, gender gaps exceed 10 percent in about half the countries. While the largest gaps favor boys, particularly in rural areas, gaps favor girls in urban and rural areas in one-third of the countries.

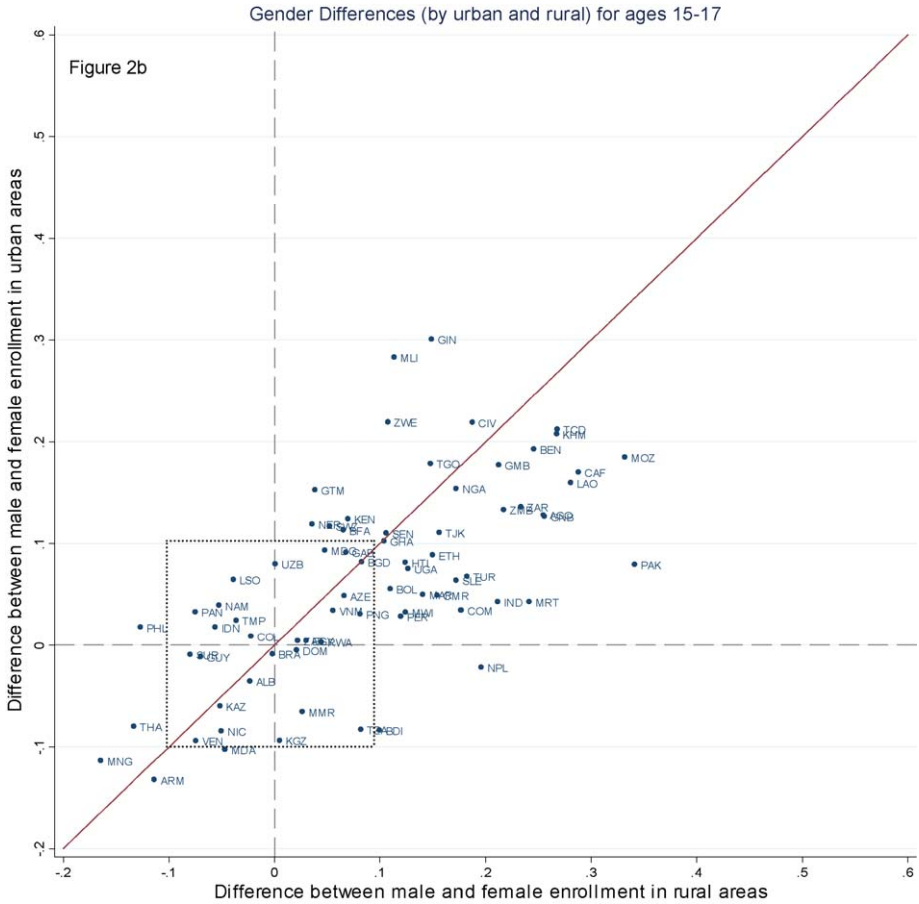

Figure 2. (*continued*)

(3) For ages 15–17, two thirds of the points lie below the 45° line, indicating larger male–female gaps in rural areas. It is in rural areas that girls' schooling mostly lags behind boys'.

(4) Girls face the greatest disadvantage in South Asian and African countries, whereas girls have higher enrollment rates than boys in both urban and rural areas in the former Soviet states which are atypically represented in the southwest quadrant.

Figure 3(a)–3(d) plots the differences in years of schooling attained by youth aged 15–24 and adults aged 25–60 for 70 developing countries. The age cutoff at 60 limits complications to our cross-country comparisons due to unequal life expectancy rates across countries. Comparing the older and younger cohorts allows us to infer changes in schooling investments across generations. In Fig. 3, panels 3a and 3b, the horizontal

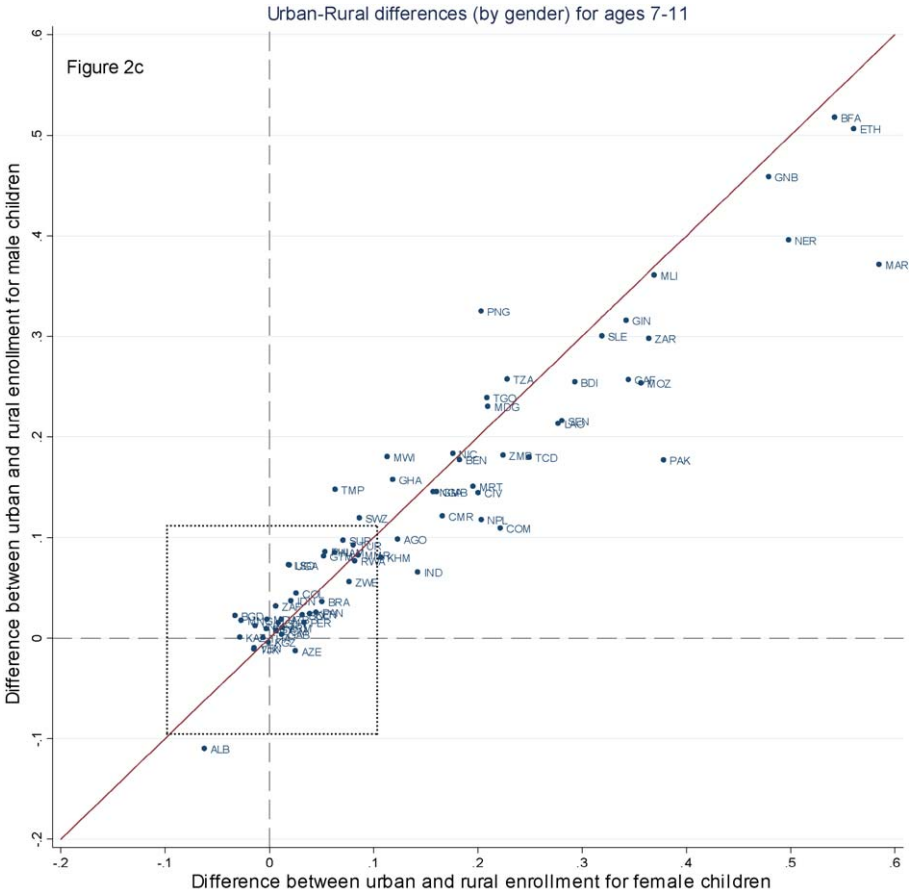

Figure 2. (*continued*)

axis shows the gender gap in years of schooling in rural areas, while the vertical axis presents the comparable indicator in urban areas. Points lying above the 45° line indicate a larger gender gap in urban areas; points below the line show larger gender gaps in rural areas. Points in the northeast quadrant imply that both gaps favor males, while points in the southwest quadrant indicate that both gaps favor females.

The graphs in 3a and 3b show several stylized facts:

(1) Women's schooling has been increasing relative to men's. In a few countries, women's gains are such that average schooling of women is now greater than that of men.

   (a) Most schooling gaps for the younger cohort are less than two years in both urban and rural areas, while most gaps for the older cohort exceed two years of schooling in urban or rural areas.

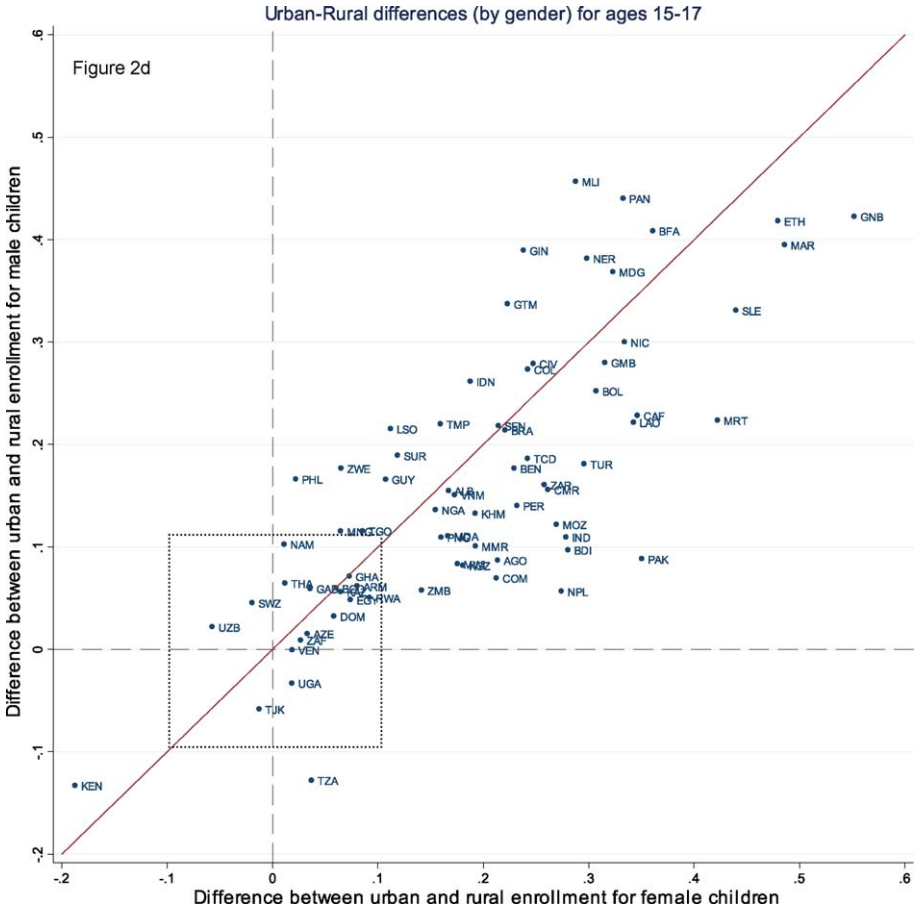

Figure 2. (*continued*)

(b) Many countries are in the southwest quadrant of the younger cohort plot but very few in the older cohort plot, indicating an increased likelihood that women attain more years of schooling than men in the younger cohorts. Correspondingly, there are fewer countries in the northeast quadrant for the younger than the older cohort, indicating a decreased probability that men's schooling exceeds women's in both urban and rural areas for the younger cohorts.

(2) In the plot of the older cohort, many countries lie above the 45° line, indicating that urban areas commonly have larger male–female gaps. In the younger-cohort plot, most countries fall below the 45° line, so the gender gaps are larger in rural areas.

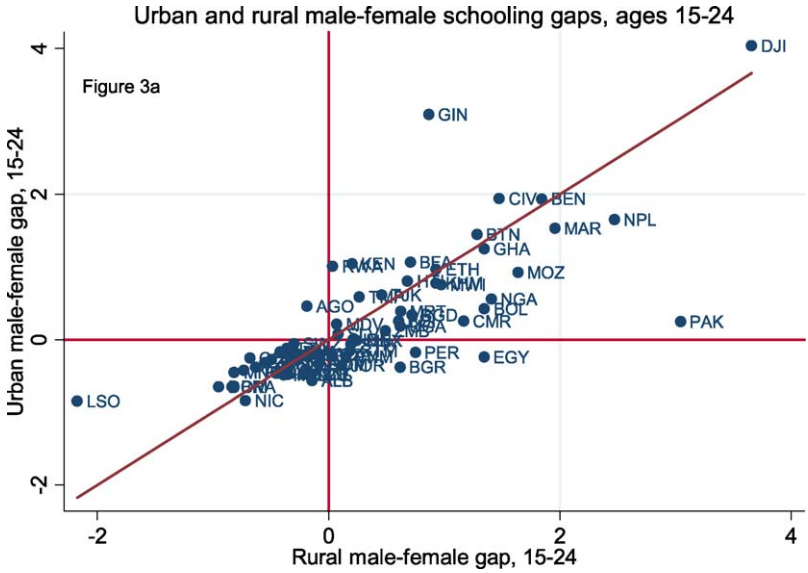

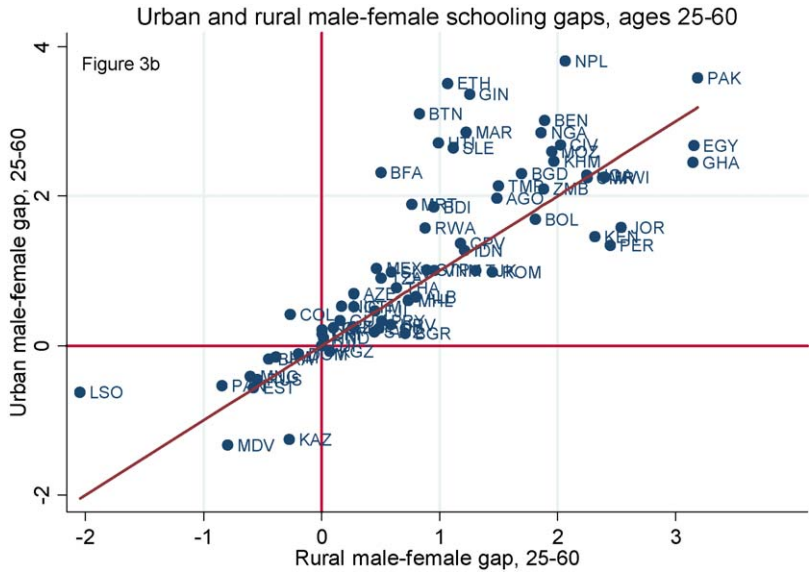

Source: Graphs computed from data from latest household surveys in 70 countries; database prepared for the World Bank, *World Development Report 2007*. All data, with few exceptions, are from year 2000 or later.

Figure 3. Urban–rural and gender gaps in years of completed schooling, ages 15–24 and 25–60.

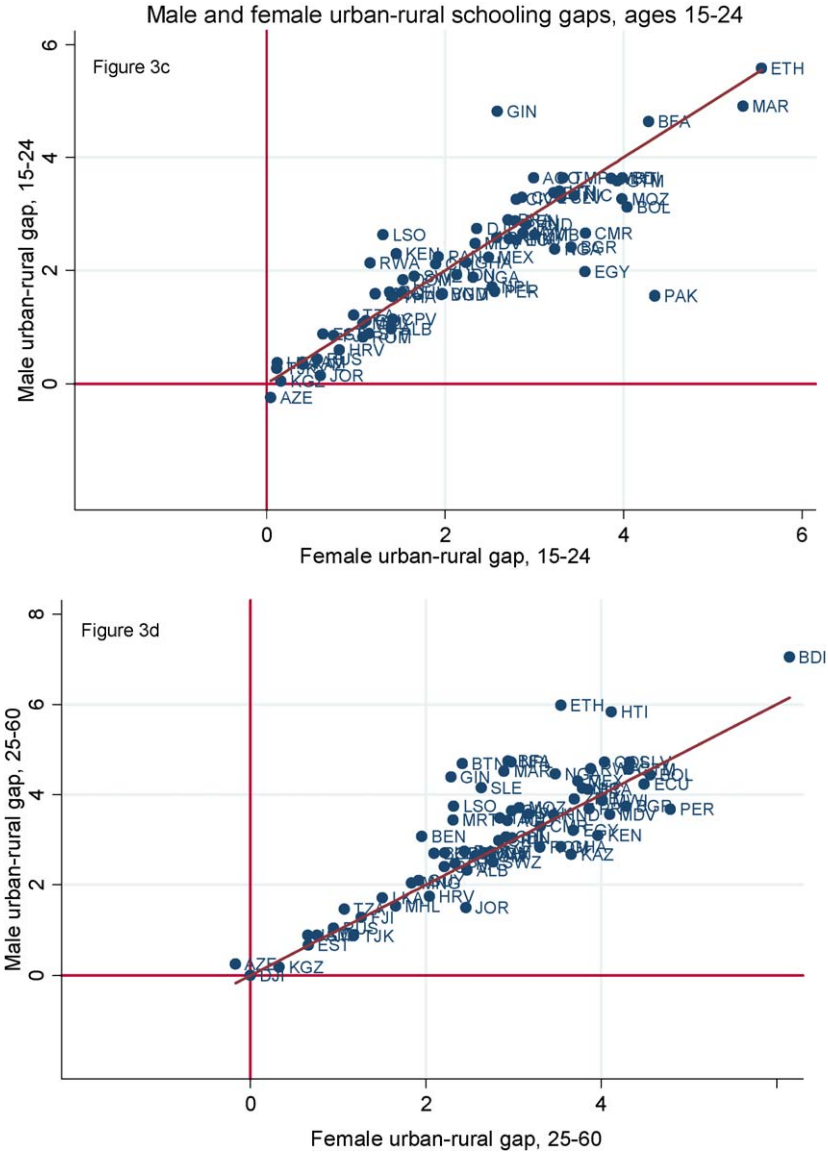

Figure 3. (*continued*)

(3) For the older cohort, the countries plotted in the extreme northeast are drawn from the Middle East, Africa and South Asia regions. Their absence from comparable positions in the younger-cohort plots suggests that it is in those areas that the most dramatic gains in women's relative to men's schooling have taken place.

Given the stopping rule for time in school implied by Eq. (2.3), it is not immediately apparent why boys and girls would spend different amounts of time in school. After all, boys and girls grow up in the same households and have equal household incomes and parent discount rates which affect schooling investment decisions. Gender differences in elements of z have ambiguous effects on boys and girls. For example, better nutrition for boys than for girls would raise the opportunity cost of schooling for boys because better-nourished boys would be more productive at work, but it would also make them more productive at school. Even improvements in the quality of schools have an ambiguous net effect because better schools imply that fewer years of schooling are needed to achieve a given level of learning.

Gender differences in work opportunities for educated labor also have an ambiguous effect because better opportunities raise both the opportunity cost of continuing in school and the potential returns to time in school. Differential returns to schooling merit some discussion here. Several recent reviews that have summarized the returns to schooling for men and women in developing countries have tended to find larger private returns to schooling for women than for men, although exceptions exist (Schultz, 1988, 2001, 2002).[5] An example of the pattern of estimated returns obtained from Mincerian earnings functions[6] is displayed in Fig. 4; the 45° line represents cases where estimated returns for men and women are identical. In only 5 of 71 cases is the return to schooling higher for men, and in 59 cases, the estimated returns were higher for women.[7] However, these estimates may be subject to measurement error and endogeneity biases; the higher average returns to schooling for women may be due to larger estimation biases in samples of women, but the existing literature lacks a systematic examination of the issue (Schultz, 1988).[8] Additional systematic investigation across countries is needed to assess whether the gap in private returns to schooling favoring women is a fiction of differential estimation bias between the sexes.

[5] Of 71 estimates each for men and women, education failed to raise earnings in only one case for women and in only 3 for men. As the consensus is that such estimates are likely lower-bound measures of true schooling returns, it seems safe to conclude that schooling generates positive private returns to both men and women.

[6] Using our earlier specification, these are estimates of the coefficient r from gender-specific regressions of the form $\ln q = \alpha_0 + rE + \alpha_x X + \varepsilon_q$, where X includes a quadratic in age, an urban dummy and marital status. Estimation is conducted separately for 71 harmonized household data sets from 48 developing countries. The data set is discussed in Fares, Montenegro and Orazem (2007).

[7] Duraisamy (2002) found higher returns for males at some education levels and for females at other education levels.

[8] Schultz (1988) finds only small differences between least squares and instrumented estimates of returns to schooling for both men and women. A more likely source of bias is differential nonrandom selection of men and women into wage work. Only 3% of women in Côte d'Ivoire and 7% in Ghana worked for wages compared to 19% and 26% of men in Côte d'Ivoire and Ghana, respectively (Schultz, 1988). It seems plausible that the more highly selected women in wage work would come atypically from the upper tail of the female ability distribution, creating a larger upward selection bias in female estimated returns. Schultz's results and those of Duraisamy (2002) suggest that selection biases for men and women are small and comparable across the sexes.

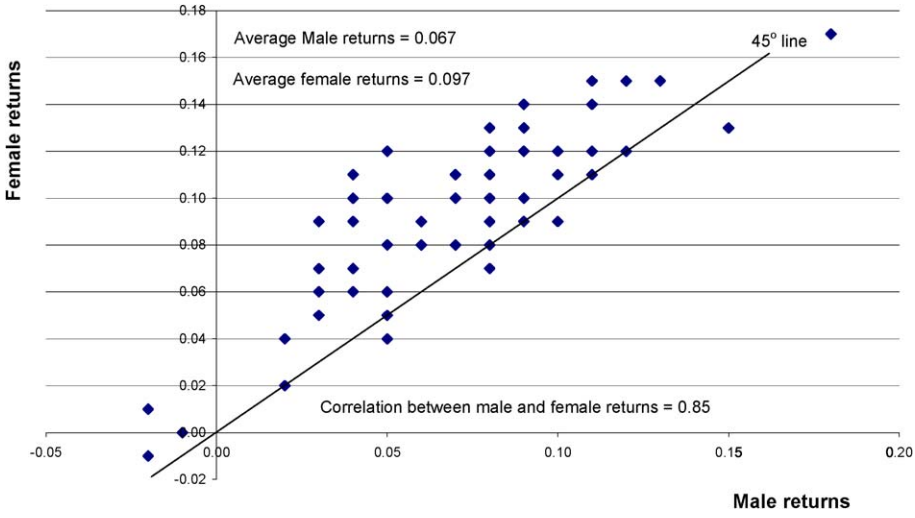

Figure 4. Paired least squares estimates of returns to schooling for males and females using household data sets from 49 developing countries, various years, 1991–2004.

Gender gaps may also reflect social norms about gender roles in familial relationships.[9] Some studies do contend that girls face higher opportunity costs of schooling due to their value in home production, although there is disagreement on how to measure the value of home time (Folbre, 2006; Hersch and Stratton, 1997; Parente, Rogerson and Wright, 2000; Smeeding and Weinberg, 2001). Indirect evidence of the impact of home production on schooling in Peru is presented by Jacoby (1993) who uses the age and sex composition of siblings (such as the presence of an older or younger sister). He finds that the number of children under five tends to raise the value of time of older children, increasing the probability of drop out.

In addition, there may be social taboos against allowing unmarried girls in public or traveling far from home but no such taboos for unmarried sons, making the cost of girls' schooling greater for girls than for boys. For example, in Pakistan special transportation or a chaperone must often be arranged for daughters in middle and secondary schools (Holmes, 2003). Social norms may also affect the returns to schooling, such as when norms specify that sons remain with their parents after marriage but that daughters move away, so parents might discount their daughters' schooling more heavily (Becker, 1985; Anderson, King and Wang, 2002; Connelly and Zheng, 2003; Quisumbing and Maluccio, 2000). If there are social taboos against allowing unmarried daughters in public or

[9] Munshi and Rosenzweig (2006), for example, examine how boys and girls respond to rapid changes in employment opportunities in urban areas in India and how caste-based networks interact with gender in determining school choice. They find that increased demand for skills can make it possible for girls from a low caste to have more schooling than boys.

being away from the family but no such taboos for unmarried sons, then the cost of girls' schooling will be greater than for boys.

In developed countries, schooling gaps between the sexes have largely disappeared as would be expected if the exogenous variables in (2.3) did not differ between boys and girls. If the source of the schooling gap in developing countries is differences in the value of child time outside of school, these differences will likely disappear as the country develops. If the source involves social norms or taboos that cause households to discount girls' schooling more heavily, the differences may persist. Also as countries impose and enforce legislative restrictions on child labor outside the home, boys would be likely to reduce their paid work while girls would continue with their responsibilities at home, thus increasing the relative cost of schooling for girls.

The public role for investing more in girls' schooling has been justified by the observed relationships between women's schooling and reduced fertility behavior, improved infant and child health, and higher cognitive attainment of children.[10] Studies have shown that mother's education improves child nutrition directly through the higher quality of care that more educated mothers can provide and through their greater ability to mitigate adverse shocks, such as food price changes, that might reduce food intake (Thomas and Strauss, 1992). In India, children of more literate mothers study nearly two hours more a day than children of illiterate mothers in similar households (Behrman et al., 1997). In Malaysia, while both the mother's and the father's education have significant positive effects on their children's schooling, the mother's education has a far greater effect than father's education on daughters' education, while the mother's and father's education have about equal, although lower, impact on sons' (Lillard and Willis, 1994). These findings underscore the gains that women's schooling can bring for improving the next generation's human capital, but are not necessarily considered by individual parents deciding how long to send their daughters to school.

2.3. Urban–rural schooling gaps

Returning to Fig. 2, panels (c) and (d) show the differences in urban–rural enrollment gaps for boys and girls of different ages. The interpretation of the urban–rural figures parallels that of the male–female plots. Points in the northeast quadrant represent countries with positive urban–rural gaps for both boys and girls, while those in the southwest represent countries with more favorable enrollment rate for rural children. Points below the 45° line imply larger urban–rural gaps for girls than boys. Among the conclusions:

(1) Urban–rural gaps are much larger than the male–female gaps. Even at the youngest ages of 7–11 years, there are many points outside the 10-percent box. The urban–rural gap exceeds 10% in half the countries. The gaps are generally of similar size for boys and girls, but the largest urban–rural gaps tend to be for girls.

[10] See Chapter 2 of King and Mason (2001) and Schultz in this Handbook for reviews of the literature in developing countries.

(2) As children age, the urban–rural gaps remain substantial with the largest gaps being for girls' enrollments. By ages 15–17, the urban–rural gap exceeds 10% in three quarters of the countries.

Figures 3(c) and 3(d) show urban–rural gaps in years of schooling for our young and old cohorts. The horizontal axis presents the urban–rural gap for females, while the vertical axis presents the urban–rural gap for males; as before, the 45° line shows combinations where the urban–rural gap is equal for males and females. Points lying above (below) the 45° line indicate a larger urban–rural gap for males (females). The plots reveal several stylized facts:

(1) Almost all countries lie in the northeast quadrant, suggesting nearly universal gaps favoring urban over rural years of schooling for both males and females. That pattern occurs for the plots for both older and younger cohorts.
(2) The range of gaps is as high as 6 years compared to about 3 in Figs. 3(a) and 3(b), and so urban–rural education gaps are larger, on average, than are male–female gaps.
(3) The countries align themselves closely to the 45° line, suggesting that urban–rural gaps are of similar size for males and females.
(4) There are more countries with gaps less than two years and fewer with gaps exceeding four years in the younger cohort plot than in the older cohort plot, suggesting that urban–rural gaps have been shrinking for some countries.

Several factors cause schooling levels in rural areas to lag behind those in urban areas. Comparisons of earnings from non-farm work between rural and urban markets generally find higher returns to schooling in urban areas (Agesa, 2001 for Kenya; de Brauw, Rozelle and Zhang, 2005 for China; and Schultz, 2004 for Mexico).[11] Indeed, Fig. 5 shows that the weight of previous empirical evidence suggests that returns to schooling are higher in urban than in rural markets. The data were derived from 66 of the household survey data sets used in Fig. 4 for which urban and rural residence was available. Results are quite consistent across countries.[12] In only three of 66 cases did schooling fail to raise earnings for both urban and rural residents of the country, although the estimated gains from schooling to rural residents are only marginally positive in ten percent

[11] Duraisamy (2002) found that returns to schooling were higher in rural areas for some education levels and higher in urban areas for others.

[12] Returns to schooling estimated from least squares may be subject to biases due to measurement error in the regressors and to unmeasured heterogeneity in ability. However, the consensus has been that these biases are of modest magnitudes (Schultz, 1988; Card, 1999; Krueger and Lindahl, 2001). However, there may be reasons why a comparison of estimated returns across urban and rural markets may yield misleading inferences. First, wages are only observed for those in wage work. Rural areas are likely to have a greater incidence of unpaid home production or self-employment, suggesting that there would be differences in the magnitude of selection bias across urban and rural markets. A second problem is that the most educated rural residents are most likely to migrate to urban markets or to be self-employed farmers, and so samples of urban workers include individuals educated in rural markets while the sample of rural workers is weighted toward the lower tail of those educated in rural schools. Schultz (1988) for Côte d'Ivoire and Ghana and Duraisamy (2002) for India found that estimated returns to schooling correcting for selection bias were very close to uncorrected measures, suggesting the bias may be modest, but the issue merits further investigation.

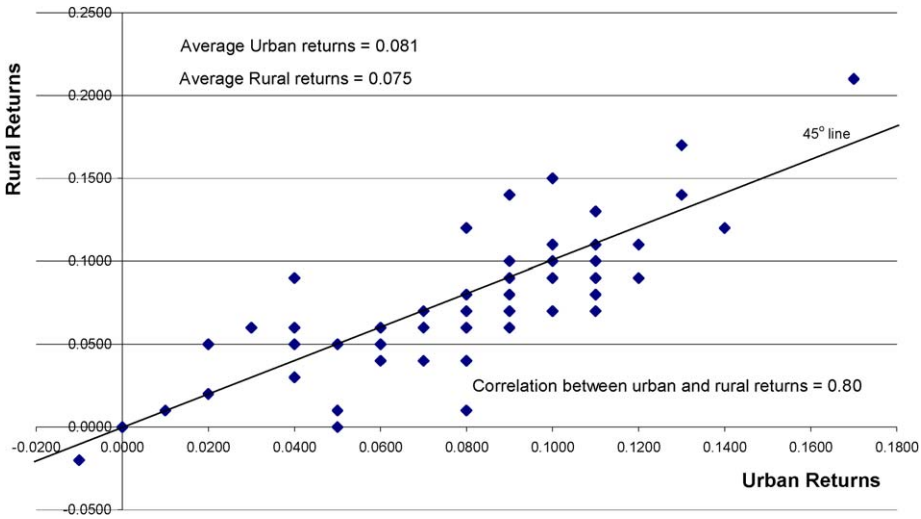

Figure 5. Paired least squares estimates of returns to schooling for urban and rural residents using household data sets from 46 developing countries, various years, 1991–2004.

of cases. Points plotted above the 45° line represents cases in which returns to schooling are higher in rural than in urban areas. In 30 of 66 pairs, the returns are equal or higher in rural than in urban areas, although urban returns are modestly higher on average.

The higher returns in urban areas provide a strong motivation for the educated to migrate from rural to urban areas. Assuming that urban and rural children have comparable latent abilities, the possibility of rural-to-urban migration means that those educated in rural areas face the same potential earnings as urban residents do. Schultz (1988) and Agesa (2001) conclude that increasing education levels in rural areas without improving employment opportunities is likely to lead to increased levels of migration: more educated youth in Latin American countries and in Kenya are likely to migrate to areas with better opportunities once they become adults.

A few other studies attempt to measure the effect of migration opportunities on schooling decisions. Kochar (2004) observes that education levels in rural areas in India are affected more by potential returns from jobs in nearby urban areas than by local wages. Schooling levels for rural males rose in areas where the wage differential between educated and less-educated male workers in urban was largest. Parents appear to pay greater attention to the employment prospects of more educated youth. Similarly, Boucher, Stark and Taylor (2007) find that rising returns to education in Mexico that can be attributed to the migration of educated rural workers to urban markets increased school attendance rates in rural areas beyond the compulsory level. In Turkey, Tansel (2002) finds that the distances to specific cities, variables that are expected to capture the effects of migration opportunities, have negative and statistically significant effects on schooling decisions, suggesting that schooling attainment is influenced by the

higher education returns expected in an urban center than in a rural area. An alternative or additional interpretation is that proximity to cities could have modernizing effects on demand for schooling. Similarly, Godoy et al. (2005) find that in a sample of male household heads from four ethnic groups in rural Bolivia, the returns to education are higher among households who live close to market towns. The enhanced returns from schooling appear to be due to off-farm opportunities as there was no effect of schooling on agricultural productivity.

That higher rural schooling may lead to rural labor migrating to urban areas is a justification for a central government role in local provision of education. To the extent that education has external benefits, this out-migration of educated workers is a net transfer of schooling returns from rural to urban areas. The loss of educated labor from rural areas may cause an underinvestment in education in those areas relative to the socially optimal level. On the other hand, to the extent that rural households are anticipating, not fearing, their children's mobility, as seems to be the case in Bolivia, India, Mexico, and Turkey, the possibility of migration could help reduce schooling inequality between urban and rural areas.

In rural areas, the return to schooling depends on the pace of technological innovation in farming and on fluctuations in farm prices. A large literature has shown that more educated farmers are the first to adopt new seeds, tillage practices, fertilizers, and animal breeds (Welch, 1970; Huffman, 1977; Besley and Case, 1993; Foster and Rosenzweig, 1996, 2004; Abdulai and Huffman, 2005). Foster and Rosenzweig (1996) examine technological growth during the Green Revolution in India in the mid-1960s and 1970s, and conclude that the return to the completion of primary school as well as the attainment of primary completion increased in areas with higher rates of exogenous technological change. Returns to schooling were greater for landowners and particularly landowning sons, supporting the hypothesis that human capital is complementary with the new technologies. Similarly, Abdulai and Huffman (2005) find that more educated farmers were the earliest adopters of crossbred cattle in Tanzania, and that the technology diffused more rapidly in areas with higher average education. Huffman and Orazem (2006) argue that human capital is critical to the process of agricultural transformation whereby improvements in agricultural productivity are sufficiently great to generate both surplus food and surplus labor needed to jump start economic growth. As farmers adopt new technologies on-farm, the resulting rise in yields result in a decline of food prices which is equivalent to an increase in real urban wages. Rising urban real wages and declining food prices create a further incentive for rural-to-urban migration of those who cannot generate sufficient returns to their skills in rural markets.

Schultz (1975) famously argued that the returns to human capital come from being better able to deal with disequilibrium. Farms that purchase inputs and sell output on the market have to respond to price fluctuations, requiring skills in finance, input and output choices, and marketing that are less needed on subsistence farms. Farmers with better skills can make better decisions regarding needed resource reallocations when rules-of-thumb are no longer appropriate. The complexity of the chemical, genetic, finance and capital investment decisions required on modern farms explains, at least in part,

why farmers and non-farmers in developed countries have more comparable schooling levels.

The return to human capital with respect to agricultural productivity is likely to be lower in farms that use traditional methods of production or where technical innovations are limited (Welch, 1970; Rosenzweig, 1980; Huffman, 1977). Hence, Schultz's (1964) observation that traditional farms are poor but efficient – historical rules-of-thumb will result in productive efficiency in environments where there are no new technological or price innovations. A large empirical literature also confirms this viewpoint. In China where it is common for household members to work in farm and non-farm activities, Yang (1997) finds that education does not enhance the labor productivity of routine farm tasks but it does increase wages from market work. Similarly, in Ghana Jolliffe (1998) concludes that returns to schooling are higher in non-farm than in farm activities, and in Pakistan Fafchamps and Quisumbing (1999) find that education has no significant effect on on-farm productivity but raises wages in non-farm work.

As a country develops, average household incomes rise, and the liquidity constraints on rural households will diminish. The agrarian sector will shrink, and labor market frictions between rural and urban markets will disappear. All of these factors will cause rural schooling attainment to rise toward the urban level, as has been observed in developed countries. For many poor countries, since income and interest rates are negatively related and since on average rural households are poorer than urban households, rural households are likely to discount the future earnings of their children more heavily. And given the preponderance of work open to children in farming activities, the opportunity cost of time in rural areas exceeds that in urban areas. These considerations result in lower rural education levels. Hence, the rationale for public subsidy of rural education is that household schooling investments based on these considerations yield less than socially optimal education levels. The external return from rural schooling may be related to the need for an educated population to react to new opportunities that arise from globalization, technological innovation, or changes in the composition of final demand for products. The rationale is even stronger if this underinvestment is due also to liquidity constraints or to parental ignorance of market opportunities for educated workers in urban markets.

3. How do government policies affect schooling gaps?

The previous section illustrates that urban–rural and male–female schooling gaps are commonly found in developing countries, and that those gaps have decreased at varying rates. This section develops a stylized model of the supply of and the demand for schooling in order to derive an equilibrium schooling investment rate first in the absence of government intervention. It then shows how alternative government policies can influence the equilibrium level of schooling. The model illustrates how differences in incomes, opportunity costs and direct costs of schooling can result in schooling gaps between urban and rural areas and how differences in the responses to those factors can

predict the effectiveness of those policies. Similarly, differences in opportunity costs can lead to differences in enrollment rates between boys and girls.

The model also shows that the measured impact of a policy on educational investment will involve numerous behavioral parameters on both the demand side and the supply side. The same policy can have different effects depending on the magnitudes of those behavioral parameters. Even this simple formulation helps to illustrate why a policy may work in some settings and not others.

3.1. A local schooling market

In this model, we assume that there is a schooling market and that, given demand for schooling, the market yields an equilibrium human capital investment rate. (We use the term "market" here to emphasize the point that demand and supply factors are at work even when schooling is completely provided by government.) For our purposes, such a market is designated by location (an urban versus a rural district) and may also be differentiated by gender (boy or girl). Let each household have one unit of child time available which would be equal to the maximum amount of educational attainment possible. Let $0 \leqslant E^D \leqslant 1$ be the average level of schooling desired by households. Let $0 \leqslant E^S \leqslant 1$ be the average share of child time for which a school space exists. The stylized demand for schooling uses our result in (2.4) that $E = E(y, p, z)$. We abstract away from the role of complementary individual or household attributes z.[13] We presume that the supply of spaces in school responds positively to its price and negatively to costs. The market is summarized by

$$\begin{aligned} &\text{Demand:} \quad E^D = \eta\big\{(1-V)p + (1-B)w\big\} + \theta(1+T)y, \\ &\text{Supply:} \quad E^S = \varepsilon p + \varphi(1-S)c. \end{aligned} \tag{3.1}$$

The demand relationship relates market household demand for schooling to three factors commonly found to influence schooling demand, including p: the logarithm of the price of schooling; w: the logarithm of the wage a child can earn while of school age; and y: the logarithm of the income level per household.[14] The parameters reflect the associated schooling demand elasticities so that $\eta < 0$ is the price elasticity of demand for schooling and $\theta > 0$ is the income elasticity of demand for schooling. The signs reflect the implications of our schooling demand formulation that household demand for schooling is stronger in local markets with low costs of attending school, low opportunity costs of child time and higher household incomes.

[13] Note that z has an ambiguous effect on years of schooling because it serves to both raises schooling productivity and increases potential earnings upon leaving school.

[14] The demand relationship used by Card (1999, Eq. (4)) Also includes a term in exogenous expected returns to schooling which is endogenous in our formulation. Adding it as an additional exogenous factor positively influencing schooling demand did not affect the conclusions from the model. The three factors included in our demand formulation correspond to the second term in Card.

We also allow three alternative public policies to influence market demand for schooling. The subsidy $0 \leqslant V \leqslant 1$ lowers the price of schooling faced by households. This can be a voucher that pays a proportion of the schooling cost or a transfer payment given directly to the school (often called a capitation grant) that is tied to a child's enrollment cost. The transfer to the household $0 \leqslant B \leqslant 1$ is made conditional on the child attending school. The transfer (referred to here as a *bolsa* following the *Bolsa Escola* programs introduced in Brazil) lowers the opportunity cost of child time in school. We contrast the conditional transfer with an unconditional income transfer $0 \leqslant T \leqslant 1$ which changes the ability to pay for schooling but does not have an explicit tie to child time use.

The market supply reflects the aggregate number of spaces provided by local public and private schools as a fraction of the population of school-age children. Market supply of schooling is assumed to depend on p: the logarithm of the price households are willing to pay for schooling; and c: the logarithm of the cost of supplying schooling services to the market. The parameters include $\varepsilon > 0$: the price elasticity of supply for schooling, and $\varphi < 0$: the cost elasticity of supply for schooling. The signs reflect standard presumptions of how aggregate supply of schooling is influenced by costs and returns. We include one more potential government intervention in the form of a cost subsidy $0 \leqslant S \leqslant 1$ which lowers the expense of providing schooling services.

The equilibrium schooling price, P^*, and equilibrium school investment rate, E^*, are

$$P^* = \frac{\varphi(1-S)c - \eta(1-B)w - \theta(1+T)y}{\eta(1-V) - \varepsilon},$$
$$E^* = \frac{\eta\varphi c(1-V)(1-S) - \varepsilon\eta(1-B)w - \varepsilon\theta(1+T)y}{\eta(1-V) - \varepsilon}. \tag{3.2}$$

3.2. Educational investments without government intervention

Virtually every government intervenes in the market for schooling in some way; in fact, school supply, especially at lower education levels, is presumed to be the purview of the government. Nevertheless, it is instructive to consider what educational investment rates would be if the government played no role in the provision of educational services so that $V = S = B = T = 0$. The equilibrium schooling price and education investments are

$$P^* = \frac{\varphi c - \eta w - \theta y}{\eta - \varepsilon},$$
$$E^* = \frac{\eta\varphi c - \varepsilon\eta w - \varepsilon\theta y}{\eta - \varepsilon}. \tag{3.3}$$

There is no guarantee that the equilibrium price is positive in the absence of government intervention. The denominator is negative, but only the first and third terms in the numerator are negative. As a consequence, private schools may not enter the market in the absence of government provision of educational services. Even with a positive equilibrium price, there is no guarantee that households will send children to school. The

equilibrium level of education investments can be written as $E^* = \varepsilon P^* + \varphi c$. The first term is positive if $P^* > 0$, but the second term is negative, and so a positive equilibrium price is a necessary but not sufficient condition to ensure positive equilibrium education investments in the government's absence.

It is useful to consider what factors increase the likelihood that $E^* > 0$ without government intervention. Markets with low cost of school provision, low opportunity costs of child time, and high household incomes are more likely to have schools, even without public support. Such conditions are more likely to exist in urban areas. Relative to urban areas, rural areas typically have lower household incomes, higher demand for child labor, and higher costs of attracting teachers and of supplying school materials. Less densely populated areas may also be unable to take advantage of returns to scale in school provision.[15] Consequently, rural areas are less likely than urban areas to have schools without some form of government intervention.[16] Girls would also have lower schooling rates if, relative to boys, the demand for their education is less income elastic and more price elastic. Such differences in elasticities would reflect differences in parental tastes for girls' schooling versus boys' schooling. [17]

3.3. Modeling different interventions

Governments typically intervene in the market for schooling. The most common rationale for public schooling provision is that there is an expected public return to schooling above and beyond the private return captured by households, and so households will under-invest in schooling compared to the social optimum. Liquidity constraints that prevent households from borrowing against future earnings may further reduce the household's choice of schooling relative to the social optimum. The most common government intervention is through the direct provision of public schools, but in many countries, the government is not able to provide enough schools to meet demand at the public school price which may not be zero but is typically much lower than cost. As a result, both developed and developing countries have experimented with other mechanisms to raise enrollment rates. We model the factors that influence the likelihood of success of these interventions within our stylized supply-demand model for schooling. While we can predict how these mechanisms might work differently in urban and rural settings, how they differ in impact on boys and girls is less clear.

[15] See Alderman, Kim and Orazem (2003) for a comparison of school costs in urban and rural Pakistan.

[16] Estimated returns to schooling are typically lower in rural areas compared to urban areas, particularly in areas characterized by traditional agriculture and low rates of rural to urban migration. If treated as exogenous, low expected returns are another factor limiting rural schooling demand.

[17] Estimated returns to schooling for women are typically higher than for men, and so exogenous anticipated returns are not likely to explain gender differences in schooling unless parents receive benefits from their boys' schooling but not from their girls' schooling.

3.3.1. Subsidies of schooling costs

In the context of the stylized model, an entirely public school system would be equivalent to setting $S = 1$, while $V = B = T = 0$ in Eq. (3.2). In this case the government fully subsidizes schooling; more generally, however, the subsidy would be partial ($0 < S < 1$) such that public schools meet additional costs by charging tuition and other fees, or public school spaces are limited and so some students have to enroll in private schools.

The effect of the school subsidy on the schooling investment is

$$\frac{\partial E^*}{\partial S} = \frac{-\eta\varphi c}{\eta - \varepsilon} > 0. \tag{3.4}$$

This means that equilibrium education rises unambiguously with the subsidy. The effectiveness of the subsidy is greatest in areas with high schooling costs, highly elastic supply responses to costs, price elastic demand for schooling and price inelastic schooling supply. Such subsidies are likely to be effective in rural areas that are characterized by high schooling costs and price elastic demand for schooling.

It is interesting to examine how government intervention in the supply of schooling services affects private provision of services. The equilibrium price falls as the government subsidy increases, as shown by

$$\frac{\partial P^*}{\partial S} = \frac{-\varphi c}{\eta - \varepsilon} < 0. \tag{3.5}$$

Consequently, if the government raises its subsidy of schooling costs by increasing direct provision of government schools with no coincident change in support for already existing private schools, it will displace some of the private school supply by lowering the equilibrium price. For example, the reduction or removal of public school fees may lead primarily to a transfer of students from private schools to public schools rather than to a net increase in enrollment. Such a crowding out effect was reported by Jimenez and Sawada (2001) in the Philippines. James (1993) found strong evidence of trade-offs between public and private school provision in developing countries.

3.3.2. Vouchers

The government may decide to allow students to use the public subsidy either in a public school or a private school. Some countries have initiated programs that give poor households vouchers that can be used to pay all or part of the tuition at a private school, while others have adopted capitation grants that transfer income directly to the school.[18] Indeed, evidence shows that private schools (though not necessarily private

[18] Evaluations of voucher or capitation grant programs in developing countries include King, Orazem and Wohlgemuth (1999) and Angrist et al. (2002) for Colombia; Kim, Alderman and Orazem (1999) for Pakistan; and Hsieh and Urquiola (2006) for Chile.

financing) is used extensively in many countries. In low-income countries, the 2003 average enrollment share of private schools was 15 percent at the primary level and 40 percent at the secondary level; in high-income countries, the corresponding shares were lower at 12 percent and 22 percent, respectively.[19]

Vouchers are often viewed as "demand-side" interventions and capitation grants as "supply-side" interventions, but they are the same policy in that they both lower the cost of schooling faced by the household. In practice, they differ in the number of transactions required to implement the policy: capitation grants require one transaction between the government and the school, while a voucher plan requires two transactions, one between the government and the household and a second between the household and the school. It is also possible that the two policies would have different effects because a voucher plan requires more involvement by the household in the transaction than a capitation grant, and thus may elicit more attention from parents on their children's schooling. In the context of our stylized model, we assume that the two policies have identical equilibrium effects on enrollment.[20]

Setting $V > 0$ and $S = B = T = 0$ in (3.2), the effect of a voucher or capitation grant on equilibrium schooling investment is

$$\frac{\partial E^*}{\partial V} = \frac{\eta\varepsilon(\varphi c - \eta w - \theta y)}{(\eta(1-V)-\varepsilon)^2} = \frac{\eta\varepsilon(\eta-\varepsilon)P^*}{(\eta(1-V)-\varepsilon)^2} \tag{3.6}$$

where P^* is the equilibrium price obtained without any government intervention in equation (3.3). The partial derivative is only positive if $P^* > 0$, or in other words, if economic conditions have induced private school entry even without government support. This implies that vouchers will not be effective in the absence of private schools, and so vouchers are more likely effective in urban rather than rural markets.

Conditional on $P^* > 0$, the effectiveness of the program in raising E^* rises as V rises. However, the second derivative is negative, and so the marginal increase in educational investment gets smaller as V increases. The schooling response to vouchers also increases as the equilibrium price of schooling increases. Consequently, it is possible that relatively modest vouchers may have positive effects on educational investment rates, even in areas with high-priced private schools.

The effectiveness of the voucher also increases as school supply responds more elastically to price. School supply is almost surely more elastic in urban areas, but will be particularly elastic in areas with excess school capacity (King, Orazem and Wohlgemuth, 1999). To the extent that additional space can be added at low cost to accommodate additional students, vouchers and conditional transfer programs will be most effective in markets characterized by excess private school capacity. Again, this favors the effectiveness of vouchers in urban rather than rural markets.

[19] These data on enrollment shares in private schools were obtained from the World Bank education database, available on http://sima.worldbank.org/edstats. Private schools pertain to church schools, private non-sectarian schools, and community schools, whether for-profit or not-for-profit.

[20] In our model, we treat the voucher as a payment to the household. It could also be treated as a payment to the school.

3.3.3. Unconditional income transfers

The relationship between household income and schooling investment has long been established. On this basis, policy analysts have suggested that income transfers or safety net programs will raise enrollment rates even without explicit conditions on school attendance. We can evaluate such a program by setting $S = V = B = 0$ in (3.2) and allowing $T > 0$.

$$\frac{\partial E^*}{\partial T} = \frac{-\varepsilon\theta y}{\eta - \varepsilon} > 0. \tag{3.7}$$

The income transfer program is most effective when schooling demand is income elastic but not price elastic and when schooling supply is price elastic so that additional space can be added without raising schooling prices rapidly. The best case for such transfers can be made as part of an income support program for low-income households that face periodic income shocks from unstable employment. Several recent studies (discussed in more detail in Section 5) have found that when poor households experience a sudden income loss due to business cycle shocks, national currency crises, or crop failures, child time is reallocated from school to work. Even temporary income shocks could cause permanent loss of potential human capital to the extent that children fall behind their peers in school and are more likely to drop out. Programs that help such households absorb transitory income shocks may allow them to keep their children in school. Nevertheless, most recently instituted income transfer programs have opted to place conditions on how child time is allocated rather than relying solely on the income effect of the transfer on child time in school.

3.3.4. Conditional income transfers

In Latin America, there has been an explosion of interventions that transfer income to poor households in exchange for a commitment to send children to school and/or to reduce child labor. Examples are Mexico's PROGRESA/*Oportunidades* program and Brazil's *Bolsa Escola* program.[21] Such a program implies setting $B > 0$ and $S = V = T = 0$. The predicted effect of a conditional transfer that exactly replaces the income previously generated by the child at work is

$$\frac{\partial E^*}{\partial B} = \frac{\eta\varepsilon w}{(\eta - \varepsilon)} > 0. \tag{3.8}$$

[21] For example, in the PROGRESA/Oportunidades program discussed by Parker, Rubalcava and Teruel in this Handbook, grants are awarded to mothers in households judged to be extremely poor every two months during the school calendar and all the children between 7 years and 18 years in these households are eligible. To receive the grant parents must enroll their children in school and ensure that children have a minimum attendance rate of 85%, monthly and annually.

The effect on schooling is unambiguously positive, reflecting a pure substitution effect toward increased schooling. More generally, for some households, the conditional transfer offers increased income while for others it will actually lower income because child time will be allocated away from labor and to school. Therefore, income is likely to change as a consequence of the conditional transfer so that $\frac{\partial y}{\partial B} \neq 0$. Incorporating the effect of the *bolsa* on income into the analysis, the impact on equilibrium education investment becomes

$$\frac{\partial E^*}{\partial B} = \frac{\eta\varepsilon w - \varepsilon\theta\frac{\partial y}{\partial B}}{(\eta - \varepsilon)}. \tag{3.9}$$

If household income rises as a consequence of the program, then $\frac{\partial y}{\partial B} > 0$ and both the substitution and income effects will raise enrollment. If $\frac{\partial y}{\partial B} < 0$, the income effect will work against the substitution effect. Nevertheless, even if household income falls in response to an income reduction, the derivative in Eq. (3.9) will be positive if the income effect is sufficiently small. Consequently, areas with income inelastic but price elastic demand for schooling can still increase enrollments through a conditional transfer, even if households lose income as a result of the program.

Conditional transfers will also be more effective in areas with high opportunity costs of child's time. This suggests that they may be particularly effective in rural areas where child labor is more prevalent. Indeed, this is the justification for Mexico's conditional transfer program which is targeted to rural areas. In addition, the program allocates larger transfers to girls than to boys on the presumption that girls' time is more valuable to the household than boys' time.[22] The factual basis for this assumption is uncertain at best because most child work is not priced by the labor market. Admittedly noisy information on actual child wages in Mexico (Schultz, 2004) and some unpublished data from Pakistan do not find large differences in market pay between boys and girls. However, if the presumption is true, then we might expect that the conditional transfers would have a larger effect on girls than on boys.

A natural question is whether unconditional or conditional transfers would raise enrollments more. An unconditional transfer has only an income effect, so its clearest advantage is in areas with price-inelastic but income-elastic schooling demand. Because the poor are likely to have a more price-elastic schooling demand, it is doubtful that unconditional income transfers would dominate conditional transfers in the low-income populations that would be targeted by such government programs. The conditional transfers also have an advantage in that there is less leakage – households only participate if they plan to meet the enrollment obligation, and these programs usually monitor this. Households that do not send their children to school receive no transfers under the rules of a conditional transfer program, in contrast to an unconditional transfer program

[22] In this program, grants at the secondary education level are higher for females, and this premium rises with the grade attended. The level of the grants was set with the aim of compensating for the opportunity cost of children's school attendance.

in which the government ends up subsidizing some households whose children are not enrolled.

The most plausible role for unconditional transfers is as a temporary income safety net that insures households from adverse income shocks. A temporary subsidy to an otherwise price inelastic household that has suffered an income shock may prove more cost effective in maintaining enrollment relative to a conditional transfer program that implies a longer-term contractual obligation.[23]

3.4. Endogenous program placement and participation

The model discussed in this section demonstrates that the impact of a government policy depends on household responses to the policy. It also shows that policies hold more promise in some settings than others. For example, school vouchers have the largest effects on educational investments in settings where private schools already exist as might be true in urban areas, while in rural areas, school construction and conditional transfers may have larger impacts on schooling demand. These decisions regarding where to locate a policy intervention and how to respond to those interventions complicate evaluation. Rosenzweig and Wolpin (1988) argue that mobile populations will migrate toward areas receiving geographically targeted benefits so that the estimated returns to the program will be subject to a selection bias: the parameters will reflect the population most likely to want the benefits rather than the true population average. Even with immobile populations, programs are likely to be placed in areas where they are expected to be most useful (Rosenzweig and Wolpin, 1986), and so the estimated impact of the program will reflect the policymakers' placement choices.

Even carefully designed pilot programs yield results that are relevant only for comparable areas and the results cannot be directly extended to dissimilar areas. If pilots are placed in areas atypically expected to prove successful, the evaluation cannot predict accurately how universal application of the policy would perform.

3.4.1. Application: Evaluating voucher programs in Latin America

Studies of the voucher programs in Chile and Colombia illustrate the difficulties in estimating the impact of government policies or programs. In 1980, Chile transferred public schools to municipalities, and teachers became municipal employees. At the same time, households were given the freedom to choose from three school types: free municipally

[23] Martinelli and Parker (2003) show that child welfare rises more from conditional than unconditional transfers when the household is bequest constrained, as would be the case for the poorest households. Households wealthy enough to be planning bequests to their children would prefer unconditional bequests. The logic follows the impact of positive income shocks in the context of the Becker–Tomes model of intergenerational transfers discussed in Section 5. Parker, Rubalcava and Teruel in this Handbook provide a general review of conditional transfer programs.

managed schools, free subsidized private schools called voucher schools, or unsubsidized private schools that charged fees. Subsidized municipal and voucher schools receive one School Subsidy Unit for every attending child. Municipal and voucher schools receive the same per-student capitation grant (Mizala and Romaguera, 2000). The subsidized schools have to admit students up to a maximum class size, while the unsubsidized private schools can select from a pool of applicants based on household and child attributes such as income or ability.

Unlike Chile's nationwide program, Colombia's PACES program which began in 1992 gave vouchers only to eligible urban youths in poor neighborhoods. Not all municipalities agreed to participate since participation required cost-sharing between the central government and the municipal government. Compared to the municipalities that did not participate, the municipalities which opted to participate were the ones that had excess capacity in their existing private schools, had apparent excess demand for available public schools, and were in better fiscal shape (King, Orazem and Wohlgemuth, 1999). Participation by municipalities, schools and households was voluntary, and so any evaluation needs to address these endogenous choices.

Evaluations of Chile's reform have yielded mixed results,[24] in part because of differences in whether or how studies control appropriately for school choices. Hsieh and Urquiola (2006) discount findings that the Chile reform has improved average test scores, repetition rates and years of schooling because the "best" public school students transferred to private schools. However, Contreras (2002) concludes that when one corrects for school choice, the evidence of improved outcomes in voucher schools becomes stronger, not weaker. Hoxby (2003) contends that studies of Chile's reform cannot yield convincing results because they rely entirely on post-program data due to the lack of baseline data pre-dating the program. Without knowing which students chose public or private schools before the program, it is hard to properly control for sorting after the program is put in place.

In Colombia, endogenous household program participation can be addressed by the fact that a few municipal governments conducted a lottery as a mechanism for allocating the vouchers for which demand exceeded supply. Angrist et al. (2002) take advantage of this lottery to identify an appropriate counterfactual group for the voucher recipients. They compare several measures of education outcomes between lottery winners and lottery losers, and find that lottery winners completed 0.12–0.16 more years of schooling, a large enough increase to raise the future annual incomes of the winners by $36–48 per year. An achievement test given to a subset of lottery participants showed higher test

[24] Examples are Mizala and Romaguera (2000) and McEwan (2001) who investigated which type of school has performed better. They approach the issue of selection bias differently. Mizala and Romaguera admitted that school choice makes a difference and found that student characteristics differ across the types of schools; they did not address this issue and so their results are biased. McEwan (2001) addressed this selection bias by using the density of schools of each school type in the municipality to identify students' school choices; by doing so, he assumed that this density variable does not belong in the achievement production function.

scores (0.2 of a standard deviation) equivalent to about one additional year of schooling. Moreover, a follow-up study (Angrist, Bettinger and Kremer, 2006) finds that the program increased secondary school completion by 15–20 percent.[25] Note, however, that these impact estimates are derived from a particular universe of students – those who lived in municipalities that had more students interested in the voucher than there were vouchers available and who met the program eligibility criteria and applied to the program and thus participated in the lottery.[26] While it seems likely that for these populations vouchers improve student performance, it does not follow that vouchers will improve student performance in other municipalities that have less interest in private schools, or that vouchers will improve the schooling of less qualified students.

There is no strong *a priori* difference in policy effects on the enrollment rates of boys versus girls, but in Colombia's program, voucher status produced higher gains for girls than for boys: a statistically significant 0.12 more years of schooling for girls, primarily through reduced grade repetition and additional time in school, and statistically significant higher test scores (0.26 higher test score), versus not statistically significant and smaller effects (0.06 more years and 0.17 higher score) for boys (Angrist et al., 2002). Since girls were not targeted by the program, these differences in impact between the sexes may depend more on demand elasticities (i.e. household tastes) than differences in public policy.[27]

Capitation grants are the supply-side variant of the voucher. A pilot project in the province of Balochistan, Pakistan subsidized the establishment of private girls' schools in ten randomly selected neighborhoods. Parents in each neighborhood were given resources to contract a school operator to open a neighborhood private school, with the level of support being tied to the number of neighborhood girls the new school could attract. This strategy was chosen over a voucher program because of the absence of preexisting schools. Kim, Alderman and Orazem (1999) compared the enrollment growth

[25] After correcting for the greater percentage of lottery winners taking college admissions tests, the program impact on test score was the same as the finding from the earlier study: test scores increased by two-tenths of a standard deviation in the distribution of potential test scores. The authors conclude that the program was very cost effective given the low cost to the government and the benefits arising from the increase in winners' earnings due to greater educational attainment.

[26] In addition, as Angrist et al. (2002) point out, voucher status and scholarship use are not deterministic: about 10 percent of the lottery winners did not use the voucher, and 24 percent of the lottery losers used some other scholarship. However, the lottery outcome could be used as an instrument for scholarship use in a 2SLS model, and when the authors do so, they find effects that are 50 percent larger than a reduced-form effect of winning the lottery. Even with this method, however, the results still apply only to the particular universe of children that were exposed to the voucher.

[27] A few studies have estimated demand elasticities with respect to school availability for boys and girls. In Ghana, using cross-section data, Lavy (1996) finds that girls' schooling is more responsive to the distances to primary and secondary school (coefficients of −0.111 and −0.020, respectively) than boys' schooling (0.009 and −0.017), although the coefficient of distance to middle school is the same (Lavy, 1996). Tansel (1997) finds essentially the same results in Ghana – that distance in the middle school has a larger deterrent effect on girls' than on boys' schooling; however, in Cote d'Ivoire, the opposite seems to be the case, at least at the primary level.

in the program neighborhoods to enrollment growth in 10 otherwise similar neighborhoods that were not chosen in the drawing. The program increased girls' enrollments by around 33 percentage points. Boys' enrollments rose as well – though at a lower rate – partly because boys were also allowed to attend the new schools, and partly because parents would not send their girls to school and not also educate their boys.[28] A similar program attempted in rural areas of Balochistan could not sustain the schools because the rural communities were too poor to generate the revenues needed to allow a school to break even (Alderman, Kim and Orazem, 2003), consistent with our discussion above that suggested vouchers would have difficulty succeeding in less densely populated markets.

4. Static model

In order to establish *ex ante* projections of the likely success or failure of alternative government interventions in the market for schooling, we need to have a model of household behavior. In Section 2, we presented a model of how long a student stays in school, based on the relative present values of the stream of costs and of expected returns to schooling. In this section, we present a model to evaluate why an intervention had the impact it did, or to project what types of households or areas would be most suitable for expanding the intervention if it proves successful. The model shows why households might invest differently in the education of daughters versus sons or why otherwise comparable households might make different allocations of child time to school in urban versus rural areas. It also shows how various government policies enter the household's decisions. Even a simple static model of the household's schooling decisions demonstrates the salient issues we wish to explore.[29]

4.1. A household model

Households are assumed to have parents, a daughter, f, and a son, m. Parents decide how to allocate their children's time between work and school, with H_f and H_m representing the proportion of time in school. Parents are assumed to derive utility from their household's consumption of goods, C; and from their daughter's and son's human capital production, q_f and q_m. The treatment of schooling as a pure consumption good ignores the potential impact of schooling on the child's future labor productivity, but such issues can only be sensibly modeled in a life-cycle setting. Nevertheless, the utility parents derive from their child's schooling can also include expected future child earnings, so the static model can easily accommodate expected returns to human capital

[28] Private primary schools in Pakistan are generally coeducational, even though government schools are more typically single sex schools. Cultural taboos against mixing the sexes are more evident as girls age.

[29] See Becker (1993) and Singh et al. (1986) for more complete models.

investments. This model ignores leisure consumption, although adding leisure into the model does not change the model's implications for schooling choices or the implied reduced form schooling demand equations.

The productivity of child time spent in human capital production depends on the availability of local technologies that improve the productivity of child time spent in school, Z. Elements of Z include the quality of local schools, teachers and curriculum; the abilities of other children in the school, and the attributes of the school management including whether it is public or private and whether it is responsive to local parents. Child learning also depends on academic ability, μ_f and μ_m, which may affect child time in school or work.

Even if the son and daughter have identical abilities and face identical school supply, parents may value the outcomes of their time in school differently. Although there may be cases where boys' education is devalued, more common are cultural prohibitions against educating girls, exposing girls to the public, or placing women in the labor market may cause parents to derive less utility from their daughter's time in school. Applying Becker's (1971) innovation of a taste for discrimination, let parents discount the utility they get from their daughter's education relative to their son's by d_f.[30] The typical case will be $0 < d_f < 1$; if parents discount boys' schooling, then $d_f > 1$, while $d_f = 1$ indicates no discounting for either child.

Incorporating these various elements into the parents' concave utility function yields

$$U = U\big(C, d_f q(H_f, \mu_f, Z), q(H_m, \mu_m, Z)\big). \tag{4.1}$$

Parents face a budget constraint that depends on the time allocations of the household members. Because the great majority of children perform household chores or work in unpaid labor for a household enterprise rather than in the formal labor market, we use the marginal product of child time in household production as the relevant opportunity cost of child time in school. Let each household member be given one unit of time. Parents allocate full time to household production activities while children can either work in the home or go to school. The concave household production function is given by $Q = Q(1, 1 - H_f, 1 - H_m; \tau)$, where τ is the technology and fixed inputs that the household has at its disposal. The proceeds of the household's production are used to purchase consumption goods at price P_C, and to purchase schooling services that are priced at P_f and P_m per unit of the girl's time and the boy's time respectively. The budget constraint is

$$Q(1, 1 - H_f, 1 - H_m; \tau) + A - P_C C - P_f H_f - P_m H_m \tag{4.2}$$

[30] The discrimination coefficient could also be rationalized as the parents' relative expected return from the girl's human capital versus the boy's human capital. d_f might also reflect parents expectations that girls are more likely than boys to leave the household as adults, or that girls are less likely to remit earnings back to the household.

where A is non-labor income from assets. Parents maximize (4.1) with respect to (4.2), and using λ as the Lagrange multiplier, the first-order conditions are

$$U_C - \lambda P_C \geqslant 0, \tag{4.3a}$$

$$d_f U_q q_{Hf} - \lambda(Q_{Hf} + P_f) \gtreqless 0, \tag{4.3b}$$

$$U_q q_{Hm} - \lambda(Q_{Hm} + P_m) \gtreqless 0, \tag{4.3c}$$

$$Q(1, 1 - H_f, 1 - H_m; \tau) + A - P_C C - P_f H_f - P_m H_m \geqslant 0. \tag{4.3d}$$

When the inequality in (4.3b) or (4.3c) is strictly greater than zero, the girl or boy spends full time in school and does no household work. Similarly, when the inequality is in the opposite direction, the girl or boy spends full time in household work. When (4.3b) or (4.3c) holds with equality, the child allocates time for both school and household work. The child is more likely to spend full time in household work when the parents' marginal utility from child time in school is low, when the child's marginal product in home production is large, and when the price of schooling is high. A girl is also more likely to devote full time to household work when her parents discount more heavily her time in school.

Assuming all first-order conditions hold with equality, the trade-off between household consumption and educational investments in a boy (or a girl) is described by

$$\frac{U_C}{U_{H_m}} = \frac{P_C}{(Q_{H_m} + P_m)}. \tag{4.4}$$

The first-order conditions have useful insights also for why educational investments and outcomes for rural children typically lag behind those of urban children. Because rural households are poorer, on average, the marginal utility of consumption will tend to be large. Child time has many productive uses in agricultural households, raising the marginal product of child time. And inequality in school provision tends to mean that school supply is lower in rural areas and/or schools are of lower quality, thus increasing the cost of schooling as well as depressing the marginal utility of child time in school. In addition, as discussed in Section 2, human capital may have greater returns off-farm than on-farm, or in urban than in rural areas, thus lowering the utility rural parents derive from their children's schooling. Taken as a whole, the ratio of relative prices of consumption to schooling will tend to be lower in rural than urban areas, while the ratio of marginal utilities of consumption relative to school time will tend to be larger in rural than in urban areas. To bring the marginal utilities in line with relative prices, rural households must raise consumption and/or reduce child time in school relative to their urban counterparts.

The first-order conditions also have useful insights for why educational investments and outcomes may differ between boys and girls. The trade-off between investments in a daughter's versus a son's schooling is described by

$$\frac{U_q q_{H_m}}{U_q q_{H_f}} = \frac{d_f(Q_{H_m} + P_m)}{(Q_{H_f} + P_f)}. \tag{4.5}$$

Absent differences in school quality or availability, differences in productivity of time in the household or in school, or differences in the price of schooling, and if $d_f = 1$, parents should invest equally in the education of their daughters and sons. If $d_f < 1$, parents will view boys schooling as less expensive, even if boy's and girl's time is equally productive in the household and if the schooling prices are identical. Parents will invest more of the boy's time in schooling, even if the two children are equally productive in school. There is considerable evidence that parents in parts of South Asia, the Middle East and Africa discount their daughters' human capital relative to their sons', particularly in rural areas. However, there may also be circumstances when $d_f > 1$, particularly in societies where educated girls generate a higher bride price that goes to the parents.[31] The tendency to invest less in girls may be further reinforced if boys and girls schools are separated, and if boys' schools are of higher quality and closer proximity.

In some settings and cultures, boys are more likely than girls to work outside the home, and in most countries, girls spend more time in household work (cooking, cleaning, child care) than boys. These time allocations do not necessarily imply that girls are favored over boys to attend school, or vice versa. It is very common for children to combine school and work in developing countries, in part because the school session lasts only a few hours per day. While child work in the home or in the market may not reduce school enrollment directly, it may affect education investments by lowering the productivity of child time in school. One might expect market work to be more damaging to child schooling than is home work. Children engaged in market work are more likely to be exposed to hazards that would cause injuries or illnesses that would disrupt school time. Work outside the home may also be more intense in terms of time or physical demands, leaving the child too tired to perform well in school.

Efforts to measure the relative impact of child home or market work on human capital production, perhaps with an interest in gauging the relative effect of girls' and boys' work on schooling attainment, face a major hurdle: almost all children do work in the home. For example, virtually all third- and fourth-grade children in a sample of Latin American countries do household chores (Table 1). There is no correlation between household work and average language test scores, whether for boys or girls or urban or rural children, but that may be due to the lack of variation in household work and not to the absence of an effect. In contrast, those children who work outside the home perform worse by 14–19 percent in language tests than those who do not work outside the home, a pattern that holds for boys and girls and for urban and rural children. These correlations do not imply causality, and the inverse relationship between test scores and market work may as easily imply that poor students are more likely to work as that child workers will be weaker students. Our point is that any effort to assess the impact of household work on schooling will fail unless there is meaningful variation in the incidence of household work.

[31] Bommier and Lambert (2000) find that in Tanzania, households send their daughters to school at a younger age and that they leave school at a younger age than their sons. They argue that the pattern of behavior is consistent with maximizing the present value of the bride price.

Table 1
Incidence of child work and average language test scores, by location and type of work

	Girls		Boys		Rural		Urban	
	No	Yes	No	Yes	No	Yes	No	Yes
Home Work Incidence (%)	3.7	97.2	5.1	95.8	4.4	96.5	4.5	96.4
Language Test Score	10.5	11.0	10.2	10.4	8.8	9.4	10.8	11.1
Market Work Incidence (%)	44.7	46.2	35.9	55.1	29.5	61.4	43.6	47.3
Language Test	12.1	10.1	11.4	9.9	10.1	8.9	12.0	10.5

Notes: Countries include Argentina, Bolivia, Brazil, Chile, Colombia, Honduras, Mexico, Paraguay, Peru, the Dominican Republic and Venezuela. Home work incidence based on child responses to the question "Do you help in household chores?" Market work incidence based on child responses to the question "Do you work outside the home?" Test scores are average number correct of 19 possible from a common test of Spanish or Portuguese skills administered in all countries in 1998.
Source: Author's calculations based on data provided by Gunnarsson, Orazem and Sanchez (2006).

Equation (4.5) suggests an avenue by which policy can equalize investments across genders if other factors conspire to induce parents to treat their boys and girls differently. If, for example, girls receive less schooling than boys, the government could lower the relative price or the relative opportunity cost of schooling for girls. The use of conditional transfers to households if they send their daughters to school will lower the opportunity cost of girls' schooling, Q_{H_f}, while vouchers would lower the direct price of schooling, P_f. Improving the relative quality or proximity of girls' schools would also tend to reduce the gender gap in schooling. Of course, if the government wished to increase parental investments in education for both children, such policies could be applied equally for both boys and girls.

4.2. Estimating the static model

Assuming an interior solution to the system of Eqs. (4.3a)–(4.3d), the reduced form equations for boy's and girl's schooling will have the functional forms:

$$\begin{aligned} H_f &= H_f(P_C, P_f, P_m, Q_{H_f}, Q_{H_m}, A, Z, \tau, d_f, \mu_f), \\ H_m &= H_m(P_C, P_f, P_m, Q_{H_f}, Q_{H_m}, A, Z, \tau, d_f, \mu_m). \end{aligned} \tag{4.6}$$

First-order approximations to these unknown functional forms are commonly used in studies of child time in school. P_m and P_f are measured by school tuition, fees and distance to the nearest school. Measures of the marginal product of child labor, Q_{H_f} and Q_{H_m}, can be estimated using a household production function or by measures of variables expected to shift the marginal product of child time but not time in school. Alternatively, researchers may use average local child wages as the value of child time, subject to the caveat that average wages are subject to considerable selection bias as most children who work do not work for wages. Furthermore, those who work may not be representative of all children who are confronted by the choice of going to school.

Measures of Z typically include the quality of the nearest school, but also the quality of household inputs that are viewed as complementary with schooling such as parental education and the availability of books or other reading materials in the home. Measures of household assets include various measures of non-labor income. Alternatively, researchers may compute the full income of the household, taken as the income that would be generated if all time were allocated solely to income generation.[32] Measures of technology typically include fixed productive assets in rural areas such as land, livestock, and farm tools and machinery. We will comment on these various measures in more detail in the next section.

The d_f is typically unobserved and will be part of the error term. However, if all the exogenous variables are available, a test of parental discriminatory preferences is to reject the null hypothesis that differences in educational investments between boys and girls can be completely explained by differences between genders in measured P_C, P_f, P_m, Q_{H_f}, Q_{H_m}, μ_f, μ_m, A, Z, and τ. Most of these measures will not differ between boys and girls, however, and those that do (presumably Q_{H_f}, Q_{H_m}, μ_f, and μ_m) are often unmeasured or unmeasurable. Of particular controversy is when the ability measures μ_f and μ_m are excluded from the analysis. Differences in observed human capital investments, H_f and H_m, or differences in observed schooling outcomes q_f and q_m may be interpreted alternatively as reflecting differences in parental or societal preferences for boys' or girls' education or as differences in underlying abilities μ_f and μ_m.[33] These interpretations invariably involve trading off potential cultural insensitivity in asserting parental favoritism with gender insensitivity in asserting differences in abilities between the sexes.

The problem of omitted variables bias is a general problem and not exclusively related to the issue of explaining differences between boys and girls. Writing the linearized reduced-form schooling demand equation for child i in household j as

$$H_{ij} = \alpha_0 + \alpha_1 A_i + \alpha_2 Z_j + \alpha_3 P_j + \alpha_4 X_j + (\gamma \mu_{ij} + \varepsilon_{ij}) \tag{4.7}$$

where the α's are unbiased regression parameters, A_i and Z_j are defined as before, P_j is a vector of child opportunity costs, and X_j is a vector of all the other household and community variables. Were ability observed, γ would be the true impact of μ_{ij} on child time in school. If instead ability were not observed, then the regression would be[34]

$$\begin{aligned} H_{ij} = {} & (\alpha_0 + \gamma\beta_0) + (\alpha_1 + \gamma\beta_1)A_i + (\alpha_2 + \gamma\beta_2)Z_j + (\alpha_3 + \gamma\beta_3)P_j \\ & + (\alpha_4 + \gamma\beta_4)X_j + (\gamma\xi_{ij} + \varepsilon_{ij}) \end{aligned} \tag{4.8}$$

where the β's are hypothetical coefficients one would obtain from an auxiliary regression of μ_{ij} on A_i, Z_j, P_j, and X_j, and ξ_{ij} is the error term in that regression. The

[32] In our application, that would be $Q(1, (1 - (H_f = 0)), (1 - (H_m = 0)); \tau) + A$.

[33] As an example, it is common for girls to outperform boys in language tests while boys outperform girls in mathematics, even when the children are in the same schools and have families with similar socioeconomic status.

[34] See Wooldridge (2002) for an excellent review of the omitted variable bias problem.

difference between the true parameter in (4.7) and the estimate derived from (4.8) depends on the magnitude of the effect of ability on observed time allocations, γ, and on the covariance between ability and the regressor, β_k.

It is plausible that the bias associated with unobserved ability, as measured by the auxiliary parameter γ in (4.8), grows with the accumulation of human capital. Parents may not know μ_{ij} until after the child has been in school long enough to allow recognition of native ability. If true, omitted variables bias should increase in magnitude the longer the child has been in school.

4.2.1. Application: The role of missing ability in evaluations of nutrition on cognitive attainment

Ability can be interpreted as the initial "school-readiness" of children. It is partially endogenous to the extent that it responds to the same variables that determine demand for schooling. Problems seen in schooling – for example, late entry, high repetition, early dropout, frequent absences, inattention and poor learning, as well as poor health of youths and adults – have been traced to malnutrition, disease, and neglect very early in the lives of children (Strauss and Thomas in this Handbook). Other studies have shown that behavioral problems in early childhood are strong predictors of high school dropout and delinquency. Children who are appropriately socialized develop the cognitive and various non-cognitive skills, including motivation, enthusiasm, cooperation, and teamwork, that allow them to develop into well-adjusted and productive adults (Carneiro and Heckman, 2003; Heckman and Rubinstein, 2001).

Several studies have examined the association between early nutritional and health status and cognitive and psychosocial skills and later educational attainment, earnings, and employment outcomes.[35] Malnutrition tends to be most common and severe during periods of greatest vulnerability, at pregnancy and the first two to three years of life. Glewwe and King (2001) conclude that malnutrition that persists into the second year of life is most critical for cognitive development, and that malnutrition in the first six months among Filipino children does not have the greatest adverse effects on child cognitive development, as argued by some observers, because it can be reversed. One estimation challenge for these studies is spurious correlation that arises due to unobserved family heterogeneity: some family or child endowment (ability) that explains a child's poor early nutritional status could explain also that child's poor cognitive development. Ignoring the ability factor attributes a causal relationship to an observed association between nutritional status and cognitive development. In principle, the problem could be

[35] For example, iodine deficiency and anemia cause poor cognitive development, particularly in children under two, and poor school performance later. In a sample of Costa Rican children, those with moderate to severe anemia had lower scores in cognitive and motor tests (Lozoff et al., 2000). See Currie (2001), Currie and Thomas (2000) and Karoly et al. (1998) for the United States; and Grantham-McGregor et al. (1997); Alderman et al. (2001), Glewwe and King (2001), Glewwe (2002), Martorell (1999), and Behrman, Cheng and Todd (2004) for developing countries. On links to adult productivities, see Strauss and Thomas (1995, 1998) for reviews.

addressed using experimental data from treatment and control groups of infants who are randomly selected, and the treatment group is provided an improved diet during the first few years of life and some years later both groups are given a test of cognitive development. One could then estimate the relationship between measures of early childhood nutrition and cognitive development, using treatment status as an instrumental variable for the latter. In the absence of an experimental intervention, Glewwe and King (2001) used community-level variables (e.g., local rainfall and prices associated with food supply) as instruments for early nutritional status, and the score in an IQ test applied before most children had entered school as a measure of cognitive achievement.

Using the same non-experimental data, Glewwe et al. (2004) examined the causal impact of early nutritional status on later schooling and academic achievement,[36] and concluded that children who were malnourished early in their lives enter school later and perform more poorly on cognitive achievement tests.[37] Using a measure of early nutritional status from longitudinal data is an improvement over using current nutritional status,[38] but it does not solve the problem of spurious correlation since parental inputs into a child's nutritional status and academic achievement are likely to be correlated and determined by unobserved (physical and cognitive) ability. To address the problem due to unobserved ability, the authors used within-family (sibling pairs) differences in IQ test scores. This controls for observed and unobserved family heterogeneity in cognitive ability, academic inputs, and nutritional inputs. The coefficient estimates of nutritional status may still be biased towards zero by measurement error so they instrument the sibling difference in heights using the older siblings' height-for-age (at birth, 12 months and 24 months) for identification.

Alderman et al. (2001) used a different approach to estimate a causal relationship between nutrition and schooling in rural Pakistan. They use the serendipitous availability of data on price shocks when the children were five years old to identify early parental inputs in child health from later inputs in schooling. Assuming that these early price shocks are uncorrelated with price shocks at later ages when schooling decisions were being made, they estimated that nutritional status is three times more important for de-

[36] Earlier studies that used cross-sectional non-experimental data found positive associations between preschool nutritional status and school achievements, but did not present persuasive evidence regarding causality. Without adequate longitudinal data, it is difficult to account for the fact that preschool nutrition reflects behavioral decisions in the presence of unobserved factors such as genetic endowments. See reviews of such studies by Pollitt (2000), Behrman (1996), and Strauss and Thomas (1995, 1998).

[37] They found that a one standard deviation increase in child height increases achievement test scores by about one-third of a standard deviation of that score or about the equivalent of spending eight more months in school.

[38] Cross-sectional data are limited in that direct measurements of preschool child health are usually not available and so this has to be inferred from current health. For example, Glewwe and Jacoby (1995) use cross-sectional data on Ghana to examine the relationship between *current* nutritional status and *current* cognitive achievement and the likelihood of delayed primary school enrollment, respectively. They find that the impact of child health on schooling is highly sensitive to the underlying behavioral assumptions and the nature of unobserved variables.

termining school enrollment than had been found in studies that treated child health as exogenous.

4.3. Estimating schooling demand

The problem of bias in estimating the true relationship between observed earnings and years of schooling has been widely studied (Card, 1999). The consensus is that omitted ability biases downward the least-squares estimate of returns to schooling, but that the bias is modest in size. There is no comparable consensus regarding the role of omitted ability on parental investments in their children's education. In part, this is due to a lack of consistency in measured child time in school, whereas there is broad consensus in the use of log earnings as the dependent variable in regressions estimating returns to schooling. The value of γ will invariably differ depending on how H_{ij} is measured.

Even if γ is nonzero, missing ability will only bias the coefficient estimates if β_k is nonzero. To the extent that parental ability is correlated with wealth and child ability is inherited from parents, household assets and child ability will be correlated, and so $\beta_1 \neq 0$. The same problem holds for parental education: if children inherit ability from their parents, then parental education will be correlated with both parents' and child's missing ability measures. For school quality, if households select schools on the basis of their children's ability, then $\beta_2 \neq 0$. This is unlikely to be a problem in rural areas where there is only one school, and so we would expect better estimates of the impact of school quality on child time in school in samples of rural children. If ability changes the value of time, then $\beta_3 \neq 0$. Evidence suggests that for young children, labor productivity is more a function of stature than cognitive development, and so β_3 is likely to be small. Community level measures are unlikely to be correlated with unobserved individual ability, and so β_4 is also likely to be small. Unfortunately, most of the policy interest is in the role of household assets and school quality on child time allocation, the variables for which the bias is likely greatest.

4.3.1. Application: Using variation in school supply to identify years of schooling completed and returns to schooling

Duflo (2001) uses a large program aimed at expanding primary school availability to generate an exogenous change in the price of schooling. The resulting change in realized years of schooling for children exposed to the program is used to identify the impact of time in school on the wages the children received in adulthood. In 1973, Indonesia used oil revenues to fund the largest primary school building project in the history of the world. Within six years, 61,807 schools were built, doubling the number of primary schools in the country and adding about 1 school per 500 school aged children. The placement of schools was aimed at meeting the greatest need as indicated by the proportion of school aged children not in school, an aim that seems to have been at least partially met.

Duflo argues that this large change in the availability of schools could be viewed as an exogenous reduction in the cost of schooling for children reaching primary school age

in 1974 (the first year that the new schools were available) relative to older cohorts. This exogenous reduction in schooling costs would shift outward the demand for schooling, and that in turn can be used to identify the impact of years of schooling on earnings. Duflo uses a 1995 sample of Indonesian men that included information on where the men were born, matching birth region with information on the number of schools built between 1973 and 1978. Men born after 1967 would have had the benefit of the school construction program for their entire primary experience. Men born before 1962 would have completed primary schooling before the building project began. Duflo compares the completed schooling and wages of the first beneficiaries (men born between 1968 and 1972) with those of the last non-beneficiaries (men born between 1957 and 1961). Because the number of schools built varied by region of the country, Duflo was able to also utilize cross-sectional comparisons of the magnitude of the effect on education and wages in regions that had large versus small increases in the number of schools.

The results are listed in Table 2. Regions are divided into two groups, high and low intensity of school construction, as measured by newly constructed schools per 1000 school-aged children. The nonrandom placement of schools clearly favors less educated areas, as the average years of schooling for the 1957–1961 birth cohort is 8.9 years in the high intensity regions and 10.4 years in the low intensity regions. Years of schooling grew faster in the high intensity school construction regions, rising by 0.5 years versus 0.4 years in the low intensity regions. Overall, the difference in education improvement in the high versus low intensity school construction regions was 0.13 years of schooling from one additional school per 1000 children. A comparable exercise relating wages to school construction finds that wages grew 3% faster for adults educated in regions with more school construction. As a check, Duflo conducts a similar exercise comparing education and wage growth between two cohorts (1952–1956 and 1957–1961), both of whom were not exposed to the school building program. The estimated changes in years of schooling and wages were much smaller and not statistically significant.

The difference in differences estimates of the impact of the school construction program on education and wages can be used to derive an estimate of the returns to schooling. Letting Z represent the intensity of school construction, H represent years of schooling and y represent the logarithm of wages, the Wald estimator for the returns to schooling would be $\frac{\partial y}{\partial H} = \frac{\mathrm{Cov}(y,Z)}{\mathrm{Cov}(H,Z)}$. In this case, a simple regression of H on Z would yield a parameter estimate $\frac{\partial H}{\partial Z} = \frac{\mathrm{Cov}(H,Z)}{\mathrm{Var}(Z)} = 0.13$ and a regression of y on Z would yield $\frac{\partial y}{\partial Z} = \frac{\mathrm{Cov}(y,Z)}{\mathrm{Var}(Z)} = 0.03$.

Missing ability means that direct estimates of $\frac{\partial y}{\partial H}$ will be biased. But because Z is an exogenous shock to both years of schooling and to earnings, the estimated derivatives of $\frac{\partial H}{\partial Z}$ and $\frac{\partial y}{\partial Z}$ are unbiased. The comparative static effect of H on y can be computed using the Wald estimator as $\frac{\partial y}{\partial H} = \frac{\partial y/\partial Z}{\partial H/\partial Z} = \frac{0.03}{0.13} = 0.23$.[39]

[39] Angrist and Krueger (1999) have an excellent review of methods to analyze labor market data including the methods reviewed in this section.

Table 2
Difference-in-differences measures of the impact of the Indonesian school building program on years of schooling and adult wages

	Years of education			Log(Wages)		
	New schools per 1000 children			New schools per 1000 children		
	Low	High	Difference	Low	High	Difference
1968–1972	10.84[a]	9.43[a]	−1.41[a]	7.48[a]	7.34[a]	−0.14[a]
1957–1961	10.44[a]	8.91[a]	−1.54[a]	7.8[a]	7.63[a]	−0.17[a]
Difference	0.40[a]	0.52[a]	0.13	−0.32[a]	−0.29[a]	0.03[b]
1957–1961	10.44[a]	8.91[a]	−1.54[a]	7.8[a]	7.63[a]	−0.17[a]
1952–1956	10.13[a]	8.56[a]	−1.58[a]	7.87[a]	7.69[a]	−0.18[a]
Difference	0.31[a]	0.35[a]	0.04	0.07[a]	0.06[a]	0.01

Source: Duflo (2001). Numbers adjusted to represent an average difference of 1 more school built per 1000 school aged children in the "High" regions versus the "Low" regions.
[a]Significant at the 0.05 level.
[b]Significant at the 0.10 level.

This is an unusually large estimate of the returns to schooling. The Wald estimate is likely influenced by other factors that affected schooling incentives that are contemporaneous with the school building program. For example, over the same time period there was at least one other major program to improve water and sanitation systems that targeted some of the same areas that received additional schools. It is plausible that improved sanitation or other contemporaneous changes may have improved child health that could also affect both schooling and wages in the same direction as the improved access to schools. To control for those factors, Duflo proposes a simplified variant of the reduced form regression (4.7) to explain years of schooling and wages (interpretable as the marginal product of schooling)

$$\begin{aligned} H_{ijk} = {} & \alpha_0 + \sum_{j=1}^{J-1} \alpha_{1j} R_{ij} + \sum_{k=1}^{K-1} \alpha_{2k} B_{ik} + \sum_{j=1}^{J} \sum_{k=1}^{K-1} \gamma_k^H Z_{ij} B_{ik} \\ & + \sum_{j=1}^{J} \sum_{k=1}^{K-1} \sum_{m=1}^{M} \delta_{mk}^H C_{ijm} B_{ik} + \varepsilon_{ijk}^H, \\ y_{ijk} = {} & \beta_0 + \sum_{j=1}^{J-1} \beta_{1j} R_{ij} + \sum_{k=1}^{K-1} \beta_{2k} B_{ik} + \sum_{j=1}^{J} \sum_{k=1}^{K-1} \gamma_k^y Z_{ij} B_{ik} \\ & + \sum_{j=1}^{J} \sum_{k=1}^{K-1} \sum_{m=1}^{M} \delta_{mk}^y C_{ijm} B_{ik} + \varepsilon_{ijk}^y \end{aligned} \tag{4.9}$$

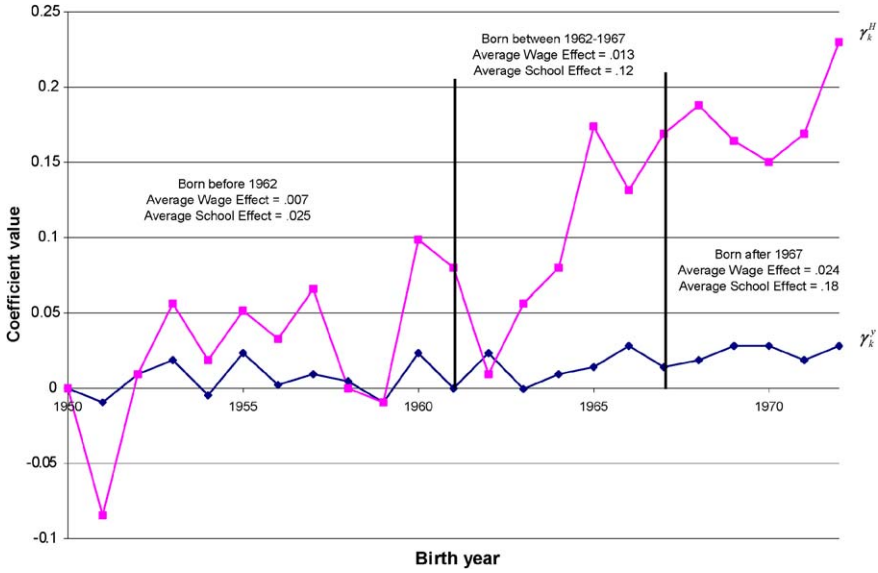

Source: Authors' compilations based on information provided in Duflo (2001).

Figure 6. Estimates of γ_k^H and γ_k^y.

for individual i in region j and birth cohort k. In this specification, birth cohort-specific effects of the regional intensity of the school building program are estimated controlling for region dummies, R_{ij}; birth cohort dummies, B_{ik}; and M region-specific variables whose impacts are allowed to vary by birth cohort. Our primary interest is in the time path of the estimates of γ_k^H and γ_k^y that are birth cohort specific estimates of the building program on years of schooling and wages.

Because of the timing of the school construction, cohorts born between 1962 and 1967 should have partial effects, and those born after 1968 should have total effects. The values of γ_k^H and γ_k^y are plotted by cohort in Fig. 6. The results are striking. Average values of the coefficients before 1962 are near zero, those in the intermediate period somewhat larger, and larger still for those born after 1967. Using the average values of the coefficients in the post 1967 period, the implied estimate of the returns to schooling is $(0.024/0.18) = 0.13$, much closer to standard estimates.

One can also use the first equation in (4.9) to identify H_{ijk} in an equation that explicitly measures the returns to schooling as in

$$y_{ijk} = \beta_0 + \sum_{j=1}^{J-1} \beta_{1j} R_{ij} + \sum_{k=1}^{K-1} \beta_{2k} B_{ik} + \beta_3 H_{ijk} + \sum_{j=1}^{J} \sum_{k=1}^{K-1} \sum_{m=1}^{M} \delta_{mk}^y C_{ijm} B_{ik} + \varepsilon_{ijk}^y. \tag{4.10}$$

Treating H_{ijk} as exogenous, the least squares estimate of β_3 is 0.077. Using Z_{ij} as an instrument for H_{ijk}, the two-stage least squares estimate of β_3 is 0.106. The finding that the OLS estimate is smaller than the 2SLS estimate is consistent with the results reported in Card's (1999) review of returns to schooling.[40]

While Duflo's estimates appear plausible, one cannot infer from the results that in general, building schools will raise enrollment or increase earnings. The schools were built where they were most needed, suggesting that the enrollment response would be atypically large. In addition, if the population can move to where the new schools are placed, the enrollment response cannot be identified by the building program as school distance becomes endogenous. The identification depends on an immobile population responding to an inflexible building rule. If the building rule flexibly responds to local demands, or if the local population changes in response to the existence or absence of the schools, then the identification breaks down.

Duflo finds that the impact of the school building project on educational attainment was greatest in less densely populated areas and in areas of greater poverty, suggesting that the price elasticity of demand for schooling is greater for the poor for rural residents. Applying a similar strategy to the case of a junior high school building project, also in Indonesia, Maliki (2005) finds that reducing the average distance to school increased school enrollment and lowered labor supply for both boys and girls, with the biggest effects in rural areas. The finding of larger impacts of school provision on rural enrollment is consistent with the implications of our stylized supply and demand model.[41]

Several recent papers have exploited the variation in school access as a plausible instrument for human capital time investments (Bedi and Gaston, 1999; Bedi and Marshall, 2002; Handa, 2002; Duflo, 2004; Foster and Rosenzweig, 2004; Glick and Sahn, 2006; and Emerson and Souza, 2006). School availability is measured alternatively by the number of schools per child or by average distance to school. School availability appears to function well as an instrument even in settings without massive new school building construction as in Indonesia, presumably because households do not relocate in response to school availability. The measure appears to be most useful in rural areas where there is more variation in school supply and where the cost of household relocation for schooling purposes is high relative to urban areas. In urban areas, variation in neighborhood school prices may be the better choice, particularly if it is difficult for households to move easily across neighborhoods.

[40] The commonly used regression specification using log earnings as a function of linear years of schooling in (4.10) to generate an estimate of the return per year of schooling β_3 is justified by the result from Section 2 that $\ln q = \ln(rV_0(E)) + rE$ where q is earnings, E is years of schooling, and $r = \beta_3$.

[41] Although these papers do not focus on private school provision, at least part of the reason that greater supply of government schools would have a smaller effect in urban areas is the possible reduction in urban private school supply. Because rural areas have fewer private schools, public provision would displace fewer rural private schools.

4.4. Discussion

Evaluations of schooling demand equations such as (4.7) should provide the key elasticities of interest for evaluating the potential values of government policies aimed at raising schooling outcomes for boys and girls in urban and rural markets. Unfortunately, despite many studies that evaluate schooling demand in different settings, conclusions regarding the relative magnitudes of these elasticities between genders and across geographic areas must be viewed as tentative. It is tempting to argue that the problem of identifying the true behavioral schooling demand parameters can be convincingly resolved only through controlled experiments. If children were randomly assigned into environments with different household incomes, schooling prices, or child wages, one could use the observed differences in human capital outcomes to derive the elasticities needed to inform policy analysis. However, there are several reasons why we may not be able to rely on social experiments alone to resolve the identification problem.

First, even the most complex social experiments discussed in this volume involve relatively small variation in the magnitude of the income transfer, school voucher, or conditional transfer that would alter the income, school price or child opportunity cost. Therefore, there is little cross sectional variation in the exogenous prices or incomes that would be needed to identify relevant elasticities.[42]

This problem is compounded by the simultaneous implementation of other social interventions in nutrition, parental education, health, sanitation, or other social services that are outside the experiment and could affect schooling decisions. As a consequence, even a rigorous evaluation that correctly identifies the overall change in schooling decisions resulting from the installation of a social program may not be able to isolate the individual impacts of income changes, versus wage changes, versus price changes on schooling demand. They may not even be able to identify the impact of the transfer from the impact of the health, sanitation, or parental education components of the intervention.[43]

Third, evaluations are best aimed at new pilot programs rather than ongoing programs (Heckman, La Londe and Smith, 1999). These pilot programs are likely to be installed

[42] This is illustrated by Jalan and Glinskaya (2003) who conclude that a large-scale school construction program in India during the period 1993–1999 had only a limited impact on enrollment. The authors compared changes in enrollment for individuals in districts that received an education program with enrollment changes of a control group of matched individuals in nonrecipient districts. Simple estimators comparing enrollment before and after the program could overstate the program's impact because enrollment increases could reflect coincident changes occurring in both program and nonprogram areas. For example, enrollment of children aged 6–10 in the program districts increased by 5.4 percentage points, only 1.3 percentage points higher than the increase in comparable non-program districts. This limited impact was attributed to the fact that the average distance to school before the program was implemented was already small.

[43] Schultz (2004) attempted to estimate an implied wage elasticity on schooling demand using schooling responses to conditional transfers in the PROGRESA/Oportunidades program. He found an elasticity of -0.2 which he concluded was unreasonably small, perhaps a consequence of the possible contamination from additional schooling responses to the various additional health and nutrition programs installed at the same time in the same places.

in areas that atypically can benefit from them. As a consequence, the responses may not reflect population averages, but only the averages for similarly suitable areas. This is not an indictment of experimental evaluation methods, only a reminder that the results of these methods cannot easily be generalized to the population as a whole. In order to generalize to other populations, one must make use of a behavioral model that can predict how dissimilar households in alternative markets will react to the same policy.

Behavioral responses to pilot programs may not be the same as responses to those same programs were they to be maintained over a long period of time and entrenched in public expectations. Non-experimental methods will still need to be used to derive many of the behavioral parameters of interest to policy makers. The most compelling non-experimental studies are ones that more closely mimic the experimental objective of randomized exposure to changes in income, wages, school prices, or school quality.[44]

One problem in summarizing the findings from studies of schooling demand is that studies do not use the same measures of schooling demand. Some (enrollment, attendance) are short-term measures while others (grade for age, years of schooling attained, measured cognitive skills) reflect decisions made over several years. We return to these measurement issues in Section 6. Nevertheless, there are sufficient empirical analyses that allow comparisons for us to advance the following conclusions about urban–rural and male–female differences in schooling:

- Schooling demand in rural areas is more income elastic and also appears to be more elastic with respect to distance to the school.
- In places where girls receive less schooling than boys (South Asia and the Middle East, rural areas of many countries), the elasticities of girls' schooling with respect to income and prices are higher than for boys.
- Income elasticities of schooling demand are generally larger in developing countries than in developed countries. Following Becker and Tomes (1986), poorer households are likely to concentrate their wealth transfers from parent to child in the form of human capital transfers. This suggests that there will be more persistence of poverty and wealth across generations in developing countries than has been observed in developed countries, a possibility that has not been subject to many empirical tests. We discuss this more completely in the next section.

These conclusions must be viewed as tentative until more systematic studies are undertaken that allow us to compare results using similar data and methods across urban and rural areas and across countries.

5. Dynamic models

Our discussion of schooling demand elasticities indicated large responses to income shocks, perhaps larger than the schooling variation caused by normally occurring cross-sectional variation in household income. Such income shocks do not fit naturally into

[44] Ravallion in this Handbook provides an extensive discussion of nonexperimental evaluation methods.

the static household production environment. To capture how unforeseen changes in prices or income affect the child schooling decision, we can embed the human capital production function into a household life-time utility maximizing model. This variant is a simplification of the model advanced by Jacoby and Skoufias (1997). The jth household is assumed to have a utility function that is time separable in its arguments. We assume that parents can only work and have one unit of time available per period. Each household has one child that also has one unit of time per period, but child time can be divided between school and work. Adult and child leisure time is assumed fixed and not subject to choice.

We assume that parents are making the decision regarding child time allocation for a schooling phase that lasts until period T. Thereafter, the child begins to work full time and to form its own household and the parents die. During each period of the schooling phase, the investment of child time produces human capital according to the concave production function $q_j^c(t+1) = q(q_j^c(t), H_j(t); q_j^P, Z_j(t))$. The child's stock of human capital depends on past acquired human capital, time spent in school in the current period, their parents' human capital endowment and an exogenous vector of local schooling inputs that complement child time in school, $Z_j(t)$. The parents' decision involves choosing a sequence of household consumption, $\{C_j(t)\}$; and a sequence of child schooling time investments, $\{H_j(t)\}$. Lifetime expected utility at time t is specified as

$$U\big(C_j(t), q_j^c(t+1)\big) + \rho_j E_t\big\{V\big(A_j(t+1)\big), \phi\big(q_j^c(T+1)\big), G(T+1), t+1\big\} \tag{5.1}$$

where ρ_j is a discount factor; $E_t\{\cdot\}$ is the expectation operator conditional on information available at time t, and $V(\cdot)$ is a value function reflecting the maximum expected lifetime utility as of the start of period $t+1$. Underlying $V(\cdot)$ is the optimal future planned sequences of $\{C_j(t)\}$ and $\{H_j(t)\}$ through period T and a planned gift of human and physical capital $G(T+1)$ that will be turned over to the child at $T+1$.

The term $A_j(t+1)$ reflects the assets available at the start of period $t+1$ if the optimal sequence of household consumption and schooling decisions were followed, and so $V(\cdot)$ can be viewed as the present value of lifetime indirect utility evaluated at time $t+1$. Because time $t+1$ decisions incorporate information unavailable at period t, we need to take the expectation conditional on information available at the start of period t in making optimal period t decisions.[45]

Assets are accumulated from income net of consumption. Household income in any period t is determined by the concave household production function $Q(\cdot)$. The intertemporal budget constraint can be written

$$A_j^*(t) - A_j(t) = Q\big(q_j^P, \big(1 - H_j(t)\big); K_j(t)\big) - P_C(t)C_j(t) \tag{5.2}$$

[45] A detailed formulation of the lifetime utility maximization model is provided by MaCurdy (1985).

where q_j^P is the human capital of the parents endowed to them by their parents; $K_j(t)$ is a vector of productive capital assets that are potentially time varying but exogenous to the household[46]; and the price of consumer goods is $P_C(t)$.

The production function presumes that better educated parents can produce more household output, but the marginal product of the child is assumed to be solely dependent on child time and is unrelated to accumulated child human capital. This assumption greatly simplifies the analysis in that it keeps the opportunity cost of child time exogenous to household decisions.[47] $A_j^*(t)$ is the value of assets at the end of period t which determines the level of assets carried into the beginning of the next period by applying the known rate of interest r. The law of motion determining asset accumulation is

$$A_j(t+1) = \big(1 + r(t+1)\big)A_j^*(t). \tag{5.3}$$

Maximizing (5.1) subject to conditions (5.2) and (5.3) and the production functions for human capital and household output we get the first-order conditions

$$\begin{aligned} &U_{C_j}(t) = \lambda_j(t)P_C(t), \\ &U_{q_j^c}(t+1)q_{H_j}^c(t) = \lambda_j(t)Q_{Hj}(t), \\ &\lambda_j(t) = \big(1 + r_j(t+1)\big)\rho_j E_t\big\{\lambda_j(t+1)\big\}. \end{aligned} \tag{5.4}$$

The first two conditions are comparable to the ones derived from the static model of household consumption and schooling investments (4.3a)–(4.3c). The first condition equates the period t marginal utility of consumption with the marginal utility of wealth sacrificed to pay for that consumption. The second condition equates the marginal utility from allocating the child's human capital to further schooling to the marginal product of allocating child time to home production.

The difference between the first two conditions in (5.4) and the earlier static formulation is the inclusion of a time condition on $\lambda_j(t)$ which summarizes all of the past and anticipated future changes to household wealth. If there were no uncertainty, $\lambda_j(t)$ would be a known sequence over the lifetime of the household, reflecting known changes in interest rates, labor productivity, prices, or any other factors that would change asset accumulation. That distinction suggests that anticipated components of non-labor asset income such as savings, investments, and debt payments are endogenous and cannot be properly included as explanatory variables in reduced form equations explaining child time use.

The third condition in (5.4) relates the marginal utility of wealth in period (t) to that in period $(t+1)$. As the other first-order conditions all relate solely to period (t), the

[46] Assets such as land, fixed buildings, or capital equipment may appreciate or depreciate at a known rate or may change productivity due to geoclimatic factors or market prices. Variable capital assets would have to be modeled jointly with decisions on time allocations.

[47] This may not be too heroic an assumption. Rosenzweig (1980) found that agricultural day wages in India did not rise with education. Alderman, Orazem and Paterno (2001) report that urban wages for children did not vary by age and presumably accumulated schooling.

third condition summarizes the dynamics of the model. To illustrate, combining the second and third conditions and substituting in for $\lambda(t)$ and $\lambda(t+1)$, we get

$$E_t\left(\frac{U_{q^c}(t+1)q_H^c(t)}{U_{q^c}(t+2)q_H^c(t+1)}\right) = \rho\big(1+r(t+1)\big)E_t\left(\frac{Q_H(t)}{Q_H(t+1)}\right) \tag{5.5}$$

where the subscript j has been suppressed. In periods where child labor productivity is relatively high, the household will optimally increase the marginal utility from time allocated to schooling, but that requires decreasing child time in school.

5.1. Application: Structural estimation of the dynamic model: Mexico's PROGRESA Program

One estimation strategy is to specify a functional form for the objective function (5.1) along with any associated restrictions and solve the lifetime maximization problem numerically. A recent example of this strategy is by Todd and Wolpin (2006) who estimate an even more complex variant of the structural dynamic model described above. Their model includes three annual decisions regarding child time use, including school, wage work, and home production. They also incorporate an additional household choice of whether to add another child to the household. They solve the dynamic programming problem for each couple in their data set through backward induction from the last period which is assumed to be the year the woman reaches age 59. The maximization involves estimating each household's expected utility from all possible sequences of fertility choices and decisions regarding child schooling, work and home time. The household is then assumed to select among these alternative time paths as in a multinomial logit formulation. The estimation also includes equations that yield the wages offered to children, the income that parents receive, and the likelihood that a child succeeds in school. Readers interested in further estimation details will have to consult the paper.

The model is applied to data on Mexican rural households. In 1998, Mexico initiated PROGRESA (now Oportunidades), a program of several simultaneous interventions aimed at improving the educational, health and nutritional status of poor rural families. PROGRESA provides cash transfers linked to youth's enrollment, regular school attendance and health clinic visits. The program also includes in-kind health benefits and nutritional supplements for children up to age five and for pregnant and lactating women. By the end of 1999, PROGRESA covered approximately 2.6 million families or about 40 percent of all rural families in Mexico. The installation of the program was randomized so that 320 villages were assigned to receive the treatment in 1998 and 186 villages were scheduled for later program installation.

Todd and Wolpin use household data collected in 1997 on all 506 villages before the program was begun. They first validate their model by examining whether it could generate within-sample enrollment and child labor rates that matched the actual distribution of child time allocations in 1997. They further validate the model by predicting how enrollment would respond to a change in the opportunity cost of schooling such

as would have occurred with the conditional transfer. The predicted changes are then compared to the actual change as reported from experimental evaluations of the PROGRESA. The model performs reasonably well for girls, predicting about an 8 percentage point increase in enrollment compared to a 10 percentage point increase from the experimental evaluations. For boys, the model predicts a 9.5 percentage point increase but the actual change was less than 5 percentage points. Note the projections only include responses to the conditional transfer. Any changes in enrollment due to the nutrition, health or other elements of the PROGRESA intervention would not be factored into the Todd–Wolpin projections and could cause their projections to deviate from the observed outcomes.

The importance of the behavioral model is that it allows Todd and Wolpin to contrast the outcomes from the conditional transfer program with other possible interventions that were not used. For example, they are able to demonstrate that unconditional income transfers would have raised enrollment by only 20% of the conditional transfer, consistent with the theoretical model that suggested the conditional transfer would have both income and substitution effects raising enrollment. They also are able to examine the sensitivity of parental schooling decisions to alternative timing and magnitude of the transfer. For example, a policy that restricts the subsidy to grades 6–9 rather than grades 3–9 but that increases the subsidy amount per child would have cost roughly the same amount but would have increased by 25% the gain in average years of schooling. Such counterfactual experiments would not be possible from the experimental evaluation alone.

It is unclear from the Todd–Wolpin exercise whether such a complex structural model is needed to generate the types of results they report. It would be useful to examine how alternative and perhaps more parsimonious behavioral models would perform in replicating the results from experimental evaluations. An additional concern is the reliance on child wages in the estimation when so few children work for wages. Only 8% of the children aged 12–15 work for wages with very few observations for the youngest children. While the authors make an accommodation for the use of such highly selected wages in their modeling, it may be that alternative strategies that would approximate the value of time in nonmarket production could generate very different results. The use of such a complicated structure makes it very costly to experiment with alternative measures of the value of time.[48]

In the end, the most promising lesson from the Todd–Wolpin study is that behavioral models can be fitted to data collected as a result of social experiments. If behavioral models can be validated by their ability to replicate the outcomes from experimental evaluations, there should be increased confidence in the counterfactual policy simulations that are based on those models.

[48] Parker, Rubalcava and Teruel in this Handbook review other examples of structural models applied to the PROGRESA case.

5.2. Dynamic labor supply equations

Interior solutions to the three first-order conditions in (5.4) result in reduced form equations for household consumption and child time in school[49]:

$$C_j(t) = C_j\big(\lambda_j(t), \overline{Q_j}(t), Q_{H_j}(t), q^c_{H_j}(t), P_C(t), q^P_j, K_j(t), Z_j(t)\big), \tag{5.6a}$$

$$H_j(t) = H_j\big(\lambda_j(t), \overline{Q_j}(t), Q_{H_j}(t), q^c_{H_j}(t), P_C(t), q^P_j, K_j(t), Z_j(t)\big). \tag{5.6b}$$

Household choices on consumption and time allocation will depend on the productivity of child time in home and school activities, on the price of consumer goods, on productive human and capital assets, and on the quality of local schools and other schooling inputs. They also depend on household full income (the amount the household could produce if all child time were allocated to home production, $\overline{Q_j}(t) = Q(q^P_j, 1; K_j(t))$) and on the marginal utility of wealth, $\lambda_j(t)$.

The sequence of $\lambda_j(t)$ is unobservable and will change over time due to new information on tastes, prices, interest rates, productivity and income. Before proceeding to the empirical specification, it is useful to illustrate how shocks to $\lambda_j(t)$ will affect $H_j(t)$. Consider a permanent, positive shock to household income caused by a permanent increase in average prices of household output. $\lambda_j(t)$ will fall because of diminishing marginal utility of wealth, as will all future $\lambda_j(t')$, $t' > t$. Rising wealth increases the household's ability to consume goods and schooling. However, $Q_{H_j}(t)$ will also increase, raising the opportunity cost of child time in school. The income and substitution effects conflict, and so the impact on child schooling is ambiguous.[50] If instead there is an increase in the productivity of adult labor in home production with no coincident change in the marginal productivity of child time, there will only be an increase in wealth and a decrease in current and future values of $\lambda_j(t)$. By the second equation in (4.4), the household would increase child time in school.

A temporary shock to household income caused by weather or price shocks that do not affect the expectation of future weather or prices will have similar but weaker predicted effects on time in school. If the household can insure itself against such transitory shocks by borrowing, saving, or taking advantage of formal or informal risk pooling mechanisms, then there may be no impact on household time allocation at all.[51]

[49] These are known as Frisch functions. See Blundell and MaCurdy (1999) for a review of the use of Frisch functions in labor supply estimation.

[50] For example, Bhalotra and Heady (2003) find an inverse relationship between land holdings and child time in school in Pakistan and Ghana, although the correlation is sensitive to covariates.

[51] Neri et al. (2006) found that households with more educated heads and higher initial incomes were best able to absorb shocks to parental employment without sacrificing their children's education. Glewwe and Hall (1998) also found that households with more educated heads were best able to maintain consumption in the face of adverse exogenous shocks related to the Peruvian business cycle.

5.2.1. Application: Jacoby and Skoufias: Enrollment responses to anticipated and unanticipated income shocks

Much of the research on household time allocation in a life-cycle setting has concentrated on household responses to permanent and transitory shocks to household income. As in many agrarian economies, rural Indian households experience unanticipated temporary increases or decreases in farm income depending on the timing and quantity of the annual monsoon rains. Jacoby and Skoufias (1997) use monthly longitudinal data from six agrarian villages from June 1975 through December 1978 to examine the extent to which child schooling can be insulated from exogenous income shocks. They further consider whether households have greater difficulty absorbing income shocks that encompass the entire community or those that hit an individual household uniquely. Their analysis is based on a variant of Eq. (5.6b). They set the child value of time equal to the local wage for market r, $W_r(t)$. They subsume all household productive attributes into the full income term $\overline{Q_{jr}}(t)$. The remaining elements are incorporated into a term representing the productivity of child time in school, $\overline{Z_{jr}}(t)$. When we first difference (5.6b), the resulting equation is of the form

$$\frac{\ln(H_j(t+1))}{\ln(H_j(t))} = \alpha_1 \Delta \ln \overline{Q_{jr}}(t) + \beta \ln W_{jr}(t) + \xi_{jr}(t) \tag{5.7}$$

where changes in the $\overline{Z_{jr}}(t)$ are assumed to make up a random error $\xi_{jr}(t)$.

Innovations in the full income term are decomposed into various common and idiosyncratic elements defined by the regression

$$\Delta \ln \overline{Q_{jr}}(t) = \gamma_{0r}(t) + \Phi'_{jr}(t-1)\gamma_1 + \left\{\Phi_{jr}(t-1) \otimes (R_r(t) - R_r)\right\}'\gamma_2 + \upsilon_{jr}(t) \tag{5.8}$$

where $\Phi_{jr}(t-1)$ is a vector of predetermined human and physical capital assets of farm j in village r; $(R_r(t) - R_r)$ is the year t deviation in rainfall from long-term village averages which is interacted with the household characteristics; and $\upsilon_{jr}(t)$ is a random error that represents remaining idiosyncratic changes in household j's full income. The Kronecker product generates a vector of interaction terms. The decomposition includes a village-specific constant $\gamma_{0r}(t)$ which provides a common unanticipated innovation to income for all households in village r. Jacoby and Skoufias argue that these aggregate shocks would be the most difficult for households to insure against. The next term, $\Phi'_{jr}\gamma_1$, is the permanent anticipated change in household full income that is forecastable based on known household attributes, which can be written as $\varepsilon^P_{jr}(t)$. The final term, $\{\Phi_{jr} \otimes (R_r(t) - R_r)\}'\gamma_2$, reflects the unanticipated idiosyncratic changes in household income due to the interaction between weather surprises and exogenous household attributes, which can be represented as $\varepsilon^T_{jr}(t)$. These changes are transitory to the extent that weather shocks do not persist across crop years.

Substituting into (5.8) yields

$$\frac{\ln(H_{jr}(t+1))}{\ln(H_{jr}(t))} = \alpha_{1r}\gamma_{0r} + \alpha_{1P}\varepsilon^{P}(t) + \alpha_{1T}\varepsilon^{T}(t) + \beta \ln W_{jr}(t) + \xi'(t+1) \tag{5.9}$$

where the various income innovation terms are allowed different parameters. If the household is fully insured against income shocks, all of these parameters will be zero. If the household cannot use formal or informal credit to smooth time allocations in response to income shocks, the coefficients will be positive and statistically significant.

There are several important issues regarding the empirical strategy for measuring permanent and transitory incomes. First, the use of full income measures is important because actual income may reflect the time allocation responses to the weather or price shocks. In that case, the measure of the transitory shock will be confounded with the consumption response to that shock. Second, it is important that the village means be taken over a long time period. Paxson (1992) argues that in short panels, it would be difficult to distinguish between permanent and transitory components of income, particularly if there is some persistence in weather levels across years. Third, the use of villages is important to allow transitory effects to be distinguished from common temporal effects across all households that confound transitory and permanent effects. As an example, a national change in tax policy could be fully anticipated and could affect all households similarly. Use of multiple villages enables such national income shocks to be controlled. Finally, the use of interactions between the rainfall shocks and household attributes allows the transitory shock to differ across households. This is the only way to test for the existence of credit relationships within villages or communities.

In the Jacoby–Skoufias sample, anticipated income fluctuations did not have much impact on child enrollments. In only one village was there substantial evidence that the anticipated aggregate shock mattered for schooling, and that was in the village with the greatest income variability over time. There was more consistent evidence that household idiosyncratic transitory shocks affected school enrollments. This was particularly true for the poorest households in the sample. The implication is that these poor agrarian households use child labor and reduced school time to help smooth consumption in the face of adverse income shocks.

If parents withdraw their children from school in the face of unexpected transitory income losses, are children permanently disadvantaged? The answer appears to be no as far as aggregate time spent in school is concerned. The simulations for India showed that even though poor households used child labor to smooth income fluctuations, the lost time in school was within 2% of the predicted school time of children in households that are perfectly insured against income fluctuations.

Much of the research on the permanent income hypothesis has concentrated on the consumption response to permanent and transitory income shocks. A common finding for farm households in the United States (Reid, 1952; Friedman, 1957; Langemeier and Patrick, 1990, 1993) and in developing countries (Wolpin, 1982; Paxson, 1993; Jacoby and Skoufias, 1998) is that consumption elasticities with respect to permanent income

are near unity and that the propensity to save out of transitory income is much larger than the propensity to save out of permanent income. These findings have led authors to conclude that farm households can insure themselves against income shocks using government safety nets, financial intermediaries, household savings or community risk pooling arrangements. And if these insurance markets immunize the farm household from transitory shocks, then there should be no labor supply response to these shocks either. Jacoby and Skoufias have shown that an important mechanism agrarian households use to attain consumption smoothing is by adjusting child time in school and work. The next section examines this question as one element of the decision of parents to transfer wealth to their children.

5.2.2. *Intergenerational wealth transfers*

Consider a household with one parent and one child. The parent can transfer wealth to the child by investing in the child's human capital and/or by giving financial bequests. Becker and Tomes (1986) provide a structure to evaluate the relative importance of these two wealth transfer mechanisms. Our simplification of their model can be viewed as a two-period variant of the dynamic optimization framework. The parents' utility function in (5.1) is reconfigured as $U(C^P, W^c)$ where the utility depends on parental consumption C^P in the first period and the children's wealth W^C in the second period. The parent can allocate own lifetime wealth, W^P, to own consumption, to human capital transfers to the child valued at G_q^P, and to financial transfers or gifts, G^P. The human capital transfers can be purchased directly through school inputs Z or by devoting time to the child

$$G_q^P = Z + sw^P q^P \tag{5.10}$$

where s is the proportion of parental time spent with children. Parental time is enhanced by the parent's stock of human capital which is valued at w^P per unit. These inputs are translated into the child's human capital by a production function $q^c = q(Z, sq^P)$. The human capital production process is subject to positive but diminishing returns and the marginal products get infinitely large as either Z or s approach zero. The child's human capital has expected value w^c.

The parent's decision is then

$$\begin{aligned}&\max_{C^P, Z, s, G^P} U\big(C^P, w^c q\big(Z, sq^P\big) + (1+r)G^P\big)\\&\quad + \lambda\big(W^P - C^P - Z - sw^P q^P - G^P\big).\end{aligned} \tag{5.11}$$

All the current period decisions are in numeraire terms to reduce unneeded prices. The interest rate on financial bequests is r. The first-order conditions are

$$\begin{aligned}&U_1 - \lambda \leqslant 0,\\&U_2 w^c q_Z - \lambda \leqslant 0,\\&U_2 w^c q^P q_s - \lambda w^P q^P \leqslant 0,\\&U_2(1+r) - \lambda \leqslant 0.\end{aligned} \tag{5.12}$$

5.2.3. Unconstrained households

If the household is unconstrained so that all first-order conditions are set equal to zero, the first-order conditions imply that

$$w^c q_Z = \frac{w^c}{w^P} q_s = (1 + r) = \frac{U_1}{U_2}.$$

The first two conditions tell us that the parent will equate the marginal productivities of expenditures on school inputs and parental time. These marginal returns to investments in child human capital are equated with the return to financial assets which is also equated to the marginal rate of substitution between parental consumption and child wealth.

The marginal returns to human capital are declining at the equilibrium while the returns to financial assets are constant at the market rate of return. Consequently, if parental wealth increases, the implied reduction in the marginal utility of income λ will result in an increase in parental consumption and in their aggregate bequest to their children. However, all of the increased bequest will be in G^P as any additional human capital transfer would have a return below $(1 + r)$.

Increases in expected returns to human capital would result in higher human capital investments. However, some of the anticipated increase in child's wealth will go toward increasing the parent's consumption. Consequently, parents can use financial instruments to smooth income fluctuations across generations. Increases in returns to financial wealth would cause parents to reduce the human capital transfer and increase the financial bequest.

5.2.4. Constrained households

If the household is liquidity constrained, as might be expected of poorer households, the last condition does not hold so that

$$(1 + r) < \frac{\lambda}{U_2} = w^c q_Z = \frac{w^c}{w^P} q_s = \frac{U_1}{U_2}. \tag{5.13}$$

As a result, the entire parental bequest will be in the form of human capital whose return will exceed $(1 + r)$. The level of investment will be inefficient in that the parent would prefer to borrow money at the market interest rate in order to finance additional schooling for the child. The child could then pay back the parent at the market rate. Constraints on such loans prevent poor households from making efficient human capital investments in their children.

Changes in the parent's income would result in an increase in the human capital transfer without any change in the financial transfer. Changes in returns to human capital investments would result in an increase in parental consumption as well as an increase in human capital transfers. On the other hand, a change in the return to financial assets would have no impact at all.

The first-order conditions for credit constrained parents show that the marginal rate of substitution between parent consumption and child wealth is higher than for unconstrained parents for whom $U_1/U_2 = (1 + r)$. This means that for parents with like preferences, the poorer (constrained) parents will devote a greater share of income to their investments in children than would their wealthier (unconstrained) counterparts.

Poor parents under-invest in their child's human capital because of the constraints on borrowing. Wealthy parents invest efficiently in their child's human capital, and that optimal level of human capital is not sensitive to marginal changes in income. Therefore, there can be a role for public transfers of human capital to poor households (through public schools or vouchers) that are funded by redistributive taxes on wealthy households.

5.2.5. *Implications for studies of intergenerational transfers for poor and wealthy households*

In wealthier, unconstrained households, marginal increases in parental income will have no impact on the child's human capital and hence on child's earnings. In poorer, liquidity constrained households, all of the transfer is in the form of human capital investments which will increase with marginal increases in parental income. Therefore, we would expect that human capital investments will be more sensitive to income changes in poor than in rich households. There will also be a larger correlation between parent and child earnings in poor than in rich households. Finally, there should be a larger correlation between parent and child wealth in poor than in wealthy households.

These hypotheses should form a strong foundation for an exploration of intergenerational transfers of human and financial capital in developing countries. As reviewed by Solon (2002) there are several studies of intergenerational wealth transfers in developed country settings with typical intergenerational wealth transfer elasticities of around 0.4. Typically these studies have not investigated the relative importance of human capital versus financial bequests, in part because they concentrate on comparisons of parent and child earnings rather than total human and financial assets. Comparable studies in developing country settings have been lacking. The model above suggests that the wealth transfer elasticities should be larger in poorer countries with more prevalent liquidity constraints than in developed countries.[52]

Another possible use of the model is to frame analyses of the link between parental income and child schooling. The model predicts that liquidity constraints should result in the largest correlations between parental income and child schooling occurring among the poorest households or countries. Indeed, estimated schooling demand elasticities with respect to household income or expenditure can be quite large in developing

[52] One study that examines intergenerational wealth transfers in developing country settings is by Dunn (2004) who finds intergenerational income elasticities of 0.69 in Brazil, much larger than those reported in Solon's (2002) review of comparable elasticities for developed countries.

countries compared to estimates for developed countries. Elasticities reported by (or derived from reported estimates) by Alderman, Orazem and Paterno (2001) and Bhalotra and Heady (2003) for Pakistan and Handa (2002) for Mozambique are near or greater than 1. The schooling response to income received as remittances from abroad also imply income elasticities greater than 1 in El Salvador (Cox-Edwards and Ureta, 2003). Edmonds, Mammen and Miller (2005) find substantial positive effects of increased cash transfers on child schooling in South Africa, with the income elasticity being greatest for the poorest households. These large income elasticities are not universal however. Glick and Sahn (2000) in Guinea, Bhalotra and Heady (2003) in Ghana, and Glewwe and Jacoby (2004) in Vietnam find income elasticities between 0 and 0.4. For perspective, Haveman and Wolfe's (1995) review of schooling demand in developed countries reports income elasticities that are all below 0.2. Their best assessment that the income elasticity is 0.1.

5.2.6. Genetic transfers of human capital

Virtually every study of schooling demand includes measures of parents' education and a measure of the full-income, non-labor income or consumption expenditures of the household. Typically, all three will have a positive and significant effect on child time in school. Some researchers have placed significance on the relative size of the coefficient on mother's versus father's schooling as reflecting the relative time each spend with the children, or their relative taste for schooling. This practice is ill-advised in that the father's education is potentially more closely tied to household income than the mother's, and so some of the impact of the father's schooling may be captured by the income variable. Even so, it is not uncommon for the coefficient on father's education to be larger than the mother's (Schultz, 2002), a result which often prompts the writer to express surprise. In fact, the transmission of parental human capital to children is undoubtedly much more complex than can be captured by the relative magnitude of reduced form coefficients.[53]

A few studies have attempted to assess if the positive linkage between parents' and children's human capital is driven by genetics or by parental attitudes toward education. Plug and Vijverberg (2003) compare the education acquired by adopted versus biological siblings to examine whether both benefit equally from the abilities of their parents. They find that the positive effect of parental IQ on their children's schooling is greater for their genetically linked children than for their adopted children. They conclude that at least 55% of the parental ability is transferred genetically.

[53] See the papers by Behrman and Rosenzweig (2002, 2005) and Antonovics and Goldberger (2005) for an example of how conclusions regarding the relative magnitude of the coefficient on mother's and father's schooling on child's schooling can vary due to seemingly innocuous decisions regarding coding years of schooling and sample inclusion. The first two papers argue that the correlation between mother's and child's schooling reflects a spurious correlation due to genetic endowments, but that father's education retains a significant positive effect even after controlling for genetic endowments. The latter concluded that there was no significant difference in the role of father's or mother's education on child education.

While the Plug and Vijverberg empirical design is plausible, an equally plausible comparison would be between the children who remain with their birth mothers and children who are placed for adoption. Because adopted children tend to be placed in households with two parents of higher than average education and wealth, it is likely that adopted children will have education levels that dominate the education received by their biological siblings who are not placed for adoption. Were such an empirical study possible, it is likely that the role of parental genetic ability (nature) would be dominated by the nurture offered by the adoptive parents. Evidence supporting this premise is provided by Zimmerman (2003) who finds that, in South Africa, fostered-in children receive the same education as children living with their own parents.[54]

5.2.7. Application: The 1998 Indonesia Currency Crisis

In January 1998, the Indonesian rupiah lost two-thirds of its value. The accompanying economic shock resulted in a 12% decline in GDP. Urban wages fell by 40% while rural wages fell by 15–20%. Such a large, widespread and apparently unforeseen shock provides an opportunity to observe how or whether different households can insulate themselves against large cyclical downturns.

Frankenberg, Smith and Thomas (2003) and Thomas et al. (2004) compare time allocation and consumption decisions of Indonesian households using two national household surveys conducted in 1997 before the crisis and in 1998 in the midst of the crisis. The methodology is largely descriptive but provides several interesting comparisons. Most important for us, they compare the responses of wealthier against poorer households, of urban against rural households, of older against younger children, and of boys against girls. Several notable results were apparent.

First, the proportional decline in per capita consumption expenditure was only half as large as the proportional decline in real wage rates. Households made up for part of the lost wages by increasing aggregate hours of work by over 40% in both rural and urban areas. Households also appeared to cope by some moving away from the heavier hit urban areas to rural areas.

Thomas et al. (2004) concentrate on how school enrollment changed across the income distribution. Wealthier households were more likely to have suffered financial losses because they owned rupiah denominated assets, while poorer households were more proportionately harmed by the decline in real wages. Note that the theory suggests that liquidity constrained households would substitute away from schooling as household incomes fall, while wealthier households may be able to borrow or consume out of assets in order to smooth consumption during the crisis. Furthermore, the decline in return to financial assets would actually induce a shift toward human capital investments

[54] Zimmerman (2003) also finds that households are more likely to accept foster children if they have no school-aged children, suggesting that the foster relationship may be based in part on household chores traded for educational access.

for wealthier households, while liquidity constrained households would be unaffected by the decline in returns to financial assets.

The pattern of enrollment changes is striking. Averaging across children aged 7–19, in both rural and urban areas, the largest enrollment declines are for households at the bottom of the distribution of per capita expenditures. For rural areas, the effect is confined to the lowest 25% of the expenditure distribution. In urban areas, the enrollment decline gradually gets smaller as household per capita expenditure rises, but only the top quintile did not experience some decline. Nevertheless, the strong prediction that enrollment of the poorest households should be most adversely affected by income changes is consistent with the Indonesian experience.

Differences in household responses to the income shock between urban and rural households are also interesting. The decrease in per capita expenditure averaged 22% in urban areas and only 6% in rural areas. However, the decline in educational expenditure per student was only 10% in urban areas but 19% in rural areas, suggesting that the larger income elasticity of schooling demand is in rural areas.

Thomas et al. (2004) also report that in poor households, enrollments declined for the youngest children most, while the enrollment of older children did not change. They argue that households are protecting the investments in their older children at the expense of younger siblings. An alternative explanation is that the poor households whose older children were still in school in 1997 were atypical of poor households generally. Most poor children aged 16–19 were already out of school before the crisis, as opposed to the near universal enrollment for children aged 8–11. That these poor households kept their children in school may reflect their atypically strong taste for schooling and not a pattern of households favoring older children over their younger siblings.

The findings that the enrollment decisions of rural households are most sensitive to cyclical shocks to household income are also found by Funkhouser (1999) for the Costa Rican recession of 1981–1983. Neri et al. (2006) found that schooling decisions of households in the lowest two income quintiles were most sensitive to job loss of the household head. None of these studies follow the children for a sufficiently long period to see if the increased child time in work during the crisis was made up later. However, aggregate enrollment statistics for Indonesia appear to have recovered by 2000, presumably reflecting both the increased ability of households to adapt over time and the gradual recovery of the Indonesian economy. There is evidence from Brazil supporting the presumption that at least some of the schooling investment lost to temporary income shocks may be regained later. Duryea et al. (2006) found that the average length of a spell of child labor in urban areas of Brazil is about four months. Combined with the evidence from Jacoby and Skoufias discussed above, it appears that transitory spells of child labor often disrupt a child's education without necessarily changing the aggregate amount of time spent in school. Nevertheless, it may be that these temporary spells out of school cause permanent loss of human capital compared to more continuous enrollment, an important question for future study.[55]

[55] Strauss and Thomas in this Handbook discuss the health consequences of income shocks.

6. Measurement matters

This section discusses the schooling input and output measurement issues one faces in estimating a static or dynamic model of schooling.

6.1. Measurement of time in school

The studies we have reviewed in this chapter have used a variety of measures of educational investments, including enrollment status, current grade or school cycle attended, and years of schooling completed. Frequently, researchers do not have a choice of which variable to use, having to rely on whatever measures are included in available survey data. However, not all measures will adequately represent the conceptual variables of interest.

Even subtle differences in the design of survey questions can yield significantly different information. The following three surveys illustrate how different series of questions can establish length of schooling:

- *Demographic and Health Survey (DHS) current core questionnaire*: Has (name) ever attended school? What is the highest level of school (name) has attended? Did (name) attend school at any time during the (xxxx) school year? During this/that school year, what level and grade [is/was] (name) attending? Did (name) attend school at any time during the previous school year, that is, (xxxx)? During that school year, what level and grade did (name) attend?
- *Living Standards Measurement Survey (LSMS)* household questionnaire (Ghana): Has (name) ever attended school? What was the highest level completed? What was the highest educational qualification attained? Did (name) attend school/college at any time during the past 12 months? How much time does (name) spend going to and from school daily? Has (name) left school now? Has (name) ever attended technical and/or vocational school? How many course-years did (name) complete? What was the highest certificate (name) achieved? Has (name) ever attended a tertiary educational institution (that is a university or college)? How many years did (name) attend? What was the last institution attended?
- *Indonesia's National Socioeconomic Household Survey, 2000 (IFLS)*, a nationally representative household budget survey conducted by the Indonesian central statistical bureau in February of each year). Of household members five years and older, the questions asked were: Is (name) in school, out of school, or has dropped out? If (name) has dropped out, what month and year did (name) drop out? What is the highest level and type of schooling (name) ever attended or is attending? What is the highest grade ever attended or being attended? What is the highest certificate earned?

It is clear from these sample questionnaires that some surveys provide data on highest grade attended but does not make clear if that grade was also completed. The LSMS and IFLS questionnaires contain questions about completion, but the DHS core questionnaire does not. The LSMS and IFLS give specific information on certificate earned

but the DHS does not. Of the three, only the IFLS contains detailed information on the timing of school leaving. Comparisons across countries are complicated by this lack of consistency in measures of schooling investment across surveys, making it more difficult to establish a better frame of reference for the magnitude and distribution of the income and price elasticities needed for policy analysis. The exchange between Behrman–Rosenzweig (2002, 2005) and Antonovics–Goldberger (2005) on whether and how parental education is transferred to children demonstrates that even when using the same data set, imprecision in questions can modify or even reverse research results.

6.1.1. Enrollment status

Many studies of child schooling use enrollment status (that is, whether or not the child is enrolled during the current or specified school year) as the dependent variable. This measure collapses the schooling decision into a dichotomous variable, and a logit or probit model is used to estimate the schooling function (e.g., Jacoby, 1993; Kim, Alderman and Orazem, 1999; Anh et al., 1998; Case, Paxson and Ableidinger, 2004). Enrollment is a noisy measure for several reasons: First, depending on when the questionnaire is applied, enrollment may or may not mean that the child will complete the current school year. Education officials in developing countries recommend taking the count of enrollees after the first month of the school year in order to get a more reliable count of the number of students for the year. Secondly, the reference period for the enrollment question may not capture information accurately. For example, Jacoby (1993) warns against using "the past 12 months" as a reference period: "The twelve month reference period may span two years, so that it is possible that the child has already dropped out of school completely by the interview date." We agree that it would be better to mention a specific school year as the reference period in surveys.

Third, being a dichotomous measure, enrollment does not reflect the actual amount of time a child spends in school; one child can be registered and attend school daily while another may also be enrolled but attend only infrequently. Consequently, children who have identical observed enrollment can have dramatically different time allocated to school. We turn to attendance measures later in the section.

Fourth, two children of the same age who are both enrolled could be in different grades because one may have started school later or may have dropped out temporarily or may have repeated a grade. Current enrollment status by itself does not capture these important differences, nor does it capture total schooling investments. Some studies attempt to deal with these limitations by defining enrollment status with respect to a particular grade level and age. For example, Anh et al. (1998) predict enrollment status for a sample of children divided into three age groups (10–12, 13–18, and 19–24); they control for age, in each age group regression as well. In Vietnam, there is a very high correlation between age and grade, so controlling for age in the enrollment regression indicates whether or not a child is enrolled in a grade as expected. They also estimate a set of logit regressions for a series of schooling outcomes (i.e., any schooling, finished primary school, has some secondary schooling, finished upper secondary school). Cen-

soring bias is a potential problem for these estimations; we turn to this issue in the next section.

Like Anh et al. (1998), Connelly and Zheng (2003) estimate a sequence of probit or logit models corresponding to enrollment in progressively higher education levels (i.e., entering primary school, graduating from primary school, entering middle school, graduating from middle school, and entering high school), with each logit equation being estimated on a sample conditioned on having finished the previous schooling cycle. By not linking the sequential decisions econometrically, however, the authors introduce a selection bias in which the sample of children who successfully complete each level may be different in ability or motivation from the group left behind.[56]

6.1.2. Years of schooling

The number of years of schooling completed is usually computed from data on the highest grade attended or completed by an individual, with enrollment in or completion of the grade being used interchangeably. Our sample questionnaires above illustrate why studies are forced to equate enrollment with completion and to ignore repeated grades and partial years of schooling completed. Specifically, unlike the LSMS or the IFLS questionnaires, the DHS core questionnaire asks only for the highest grade attended, not highest grade completed. The presumption is that the highest grade attended was, in fact, completed. If a student dropped out in the middle of the year, say, after the survey was conducted, then the highest grade completed is really one year less than the highest grade attended. This is not to say that a partial year of attendance, though not reflected in the reported highest grade completed, did not contribute at all to the accumulation of human capital.[57]

Questions eliciting information on the current grade attended will overestimate years of schooling attained for children who drop out before the year ends. For students still enrolled, current grade attended will underestimate years of schooling if they plan to continue in school. In fact, even youths who are not currently enrolled could return to school the following year, so their currently observed highest grade attended could potentially underestimate final schooling attained.[58] These considerations imply that

[56] If μ_{ij} is a common element across several equations, it is possible that its influence can be controlled through joint estimation of the equations. This requires imposing some structure on the problem, either exclusion restrictions on some equations so that the same regressors do not appear in every equation or to exploit nonlinearities in the regression equations due to the use of limited dependent variables as in Bommier and Lambert (2000).

[57] In some cases, even this is not available. Some surveys ask a respondent whether he or she has ever attended the primary or secondary cycle, so the discrete measure of schooling pertains to a range of years rather than a single year.

[58] In countries where grade repetition is very high, highest grade completed significantly underestimates the level of investment made by the student, the family, and the community. Most survey data, however, do not contain information on grade repetition.

the number of years of schooling is best used as a measure when the study population is older, such that further schooling is unlikely. However, excluding the younger population in predicting demand for schooling could miss some important changes in household decision-making and policy. Indeed, in settings where educational development is progressing rapidly, focusing only on the completed years of schooling of adults will fail to capture the dynamics of the rising educational attainment of youths. If school-age children are included in an analysis of schooling outcomes, then ordinary least squares is an inadequate estimation method and empirical methods that can utilize censored data on years of schooling will need to be considered.

Various empirical approaches have been proposed to deal with these problems. One estimation strategy that addresses these issues is the censored ordered probit model used by Lillard and King (1984), King and Lillard (1987), Glewwe and Jacoby (1995), Lavy (1996), Behrman et al. (1997), Bommier and Lambert (2000), Glick and Sahn (2000) and Holmes (2003). This approach explicitly takes into account that the reported years of schooling is the outcome of an ordered series of discrete choices about enrollment and completion. By combining information on a child's enrollment status, this method allows for enrollees to be treated differently than the non-enrollees and thus is able to differentiate between highest grade attended and highest grade completed. This is accomplished by specifying different likelihood functions for the two groups which makes it possible to estimate unbiased estimates of predicted completed schooling for a given child.

There is an implied assumption that older children who have never attended school will not enroll in the future. Data from many developing countries belie this assumption. Although the official entry age for the primary grades in most countries is six or seven years, the age-enrollment profiles of the four countries in Fig. 1 show that in rural Ethiopia and Tanzania, many children enter school for the first time at age twelve or thirteen.[59] In these settings, researchers will likely incorrectly treat late entrants as never enrolling and thus underestimate the years of education completed by children who enter later. Researchers have attempted to address this by limiting their samples to children closer to the age when most enter school. For Malaysia and the Philippines, on the basis of age of entry data, King and Lillard (1987) used eight years as the cut-off age for their sample; for Guinea-Conakry, Glick and Sahn (2000) used nine years as the cut-off age; for Vietnam, Anh et al. (1998) used 10 years. Clearly, there is a trade-off between trimming off too many young children from the sample and thus losing the information from those for whom the observed schooling data are legitimate, and attributing zero years of schooling to those who will eventually enter school three or four years late.

6.1.3. Grade-for-age

A simple grade-age ratio, used during early attempts to eliminate the (right) censoring bias discussed above, is predicated on an expected or "official grade-for-age" and does

[59] Bommier and Lambert (2000) observe that a child has only a 50 percent probability of entering school before age nine in their Tanzania sample of 5000 households.

not recognize the cumulative nature of falling behind as a child ages, such as we have seen in the age-enrollment pyramids in Section 2. An alternative measure that explicitly uses median schooling levels as a benchmark for assessing individual schooling attainment has been proposed by Joshi and Schultz (2005). This normalized Z-score measure of years of education completed is similar to that in the nutrition field. In nutrition, the weight-for-height Z-score is used as a measure of long-term malnutrition, and height-for-age Z-scores as a measure of stunting. In like manner, a schooling Z-score measures the deviation from the median years of schooling for children of that age (and sex) in the population, and this difference is then divided by the standard deviation of the years of schooling of the children of the same age (and sex). This strategy recognizes the growing dispersion in years of schooling completed as a cohort ages, and converts all the deviations to more equivalent units of standard deviations.[60] It collapses the information about an individual child's schooling relative to a reference population into one number, providing a useful measure for interpreting censored information.

Not surprisingly, the usefulness of this measure depends on the choice of the reference population. For single-country analysis, should one choose a single reference population and thus be able to compare how any child in one part of the country does relative to others? Or should one use separate reference populations by urban–rural residence or by gender, in order to have a better fit between the censored schooling attainment and the reference grade-for-age curve? Again returning to Fig. 1, we see that there are clear differences between the urban and rural populations and between the male and female populations in at least three of the countries, and so using four reference populations for each country, instead of one, to compute the Z-scores would lead to different interpretations of schooling levels. For example, borrowing the language of the nutrition literature, the Z-scores might suggest that a rural girl is not "educationally stunted" relative to other rural girls but is severely "stunted" relative to urban girls or urban boys. It would also be difficult to compare the elasticities of policy variables computed from the relationships estimated using these Z-scores.

This method also does not lend itself easily to comparisons across time and across countries. While in nutrition the height-for-age curve can be expected to remain stable for at least a few years once a reference population has been chosen, the same is not true for the grade-for-age curve. In fact, given the rapid changes in educational attainment that might follow, say, a large expansion of school supply, it is possible to observe notable shifts in the reference grade-for-age curve within just two to three years. This would be likely in areas where a large proportion of children previously never attended school. During a period of school construction, observed changes in the Z-score could be due either to changes in an individual child's schooling or to shifts in the distribution of the reference population (in terms of its mean or variance), or to both, and the

[60] Even with this Z-score measure there may still be some systematic variations in outcomes by the child's age (arising in part from other groupings), but this can be dealt with by including the child's age in years as an additional control variable in the regressions.

estimated model of demand for schooling would not be able to distinguish among these sources of change.

Similarly, comparisons across countries using this measure are difficult because one has to choose one reference population in order to be able to interpret the Z-scores. In the nutrition field, cross-country comparisons are made easier by the international community agreeing to use a single reference population. In schooling, the choice is complicated by the fact that countries have different official age of entries and enforcement, and different lengths of the education cycles, and so different "natural" grade-for-age curves.

6.1.4. Attendance: Days or hours of school

School attendance or absence during the year adds another dimension to time spent in school which might be termed "intensity." Unfortunately, answers to the survey question "Does your child attend school?" are more likely to reflect enrollment decisions and not how regularly the child actually attends. Surveys that elicit how often the child attended school the previous week are better but will still measure annual attendance with error. Parents' impressions of their child's attendance record are likely fraught with error, and students' responses to those questions may be even less reliable. Memory lapses are one reason; another is that parents and students may not be aware of which absences from school are authorized free days and which are not. The latter is not surprising when teachers themselves take many unofficial leave days.[61]

Bedi and Marshall (2002) model schooling decisions one day at a time, but takes as given that a child is enrolled for the school year. Daily attendance is computed by subtracting the number of days a child is absent from the number of days the school was open during the school year, and this estimate is used as an input in a production function of student achievement. Information about attendance is taken from school administrative records. It would have been better to combine official attendance records from the school with household survey data, but this has not been done frequently in practice. An exception is King et al. (2002) who compare school attendance records with household survey data and find that spot checks of student attendance corroborate official school records, even in schools where spot checks of teacher attendance are at considerable variance with the teacher attendance registry. Perhaps there is less incentive for teachers to misreport student attendance than their own attendance.

School attendance or absences can be a critical measure of child's actual exposure to schooling especially in settings where child work is common. In these contexts, whether a child is enrolled or not is a poor signal of how much time a child spends in school since work and schooling are not likely to be mutually exclusive (Ravallion and Wodon, 2000; Patrinos and Psacharopoulos, 1997; Gunnarsson et al., 2006; Rosati and Rossi, 2003;

[61] A study of teacher absences in six developing countries found high teacher absenteeism – 19 percent of days, on average, of which only approximately one-tenth were authorized leave days (Chaudhury et al., 2006).

and the review by Edmonds in this Handbook). In places where the school day may be as short as four hours, the time constraint does not prevent the child from being enrolled while also working. For example, Angrist et al. (2002) observe that Colombia's voucher program led to a decrease in child labor for the beneficiaries, as well as to an increase in years attained and test scores. Some children work too many hours which can decrease their school attendance, limit their ability to do homework, and lower their school performance. Longer school days and school terms may influence a child's learning, as well as reduce child labor. Yap et al. (2006) find that Brazil's after-school program in rural areas greatly reduced child labor. Differences in the length of the school term between black and white schools in the United States in the segregated era explain differences in school achievement (Orazem, 1987) and earnings (Card and Krueger, 1992) between blacks and whites.

In general, the length of the school term and the number of school hours each day are often standardized within countries, so they may not prove useful in explaining variation in academic success within a single country. Exceptions are when shortages of school spaces justify multiple shifts during the school day which shorten the school day for some students. In decentralized systems, local governments tend to have the authority to specify these durations for their schools as long as the minimum set by the central authority is met.

6.2. *Measurement of school outcomes*

Time spent in school is an input into the educational production process and not a measure of schooling outcomes. Almost any public policy related to education aims to raise human capital and not just enrollment or attendance. In fact, the few studies that have included both measures of cognitive achievement and years of schooling find it is the former that best explains variation in adult earnings (Glewwe, 2002). Therefore, it is extremely important to understand how to produce literacy and numeracy most efficiently in order to foster human capital development in poor countries. To that end, many researchers have estimated equations explaining schooling outcomes as a function of school attributes that are special cases of

$$q_k = q\big(Z_k, \bar{Z}, H_k, A_k, \mu_k\big) \tag{6.1}$$

where q_k is a measure of average cognitive attainment in school k, Z_k is a vector of variable school inputs $(Z_{1k}, Z_{2k}, \ldots, Z_{nk})$ that could include teaching personnel, materials, equipment, and methods; $\bar{Z}$ are fixed inputs that cannot be influenced by the school manager; H_k is a measure of the intensity of child time allocations to the school; A_k is a measure of the socioeconomic status of the parents; and μ_k is a measure of academic ability. Consistent with equation (4.6), parents will decide how much child time to devote to the school so that $H_k = H(Z_k, \bar{Z}, A_k, \mu_k, W_k)$ where the last term is a vector of other factors. Suppose that the school administrator decides on the mix of school inputs Z_k so as to maximize child attainment (6.1), given knowledge of child abilities and community wealth, given the parent's child time allocation decision, and

given a budget constraint $B_k = \sum_n p_i Z_{ik}$ with attached Lagrange multiplier λ. The school's first-order conditions will be

$$\frac{\partial q}{\partial H_k}\frac{\partial H_k}{\partial Z_{ik}} + \frac{\partial q}{\partial Z_{ik}} - \lambda p_i \leqslant 0. \tag{6.2}$$

The school's reduced form input allocation decisions will be of the form $Z_{ik} = Z_i(B, p_1, p_2, \ldots, p_n, \bar{Z}, A_k, \mu_k)$. Schools will decide on the mix of inputs incorporating how those decisions will affect parents' decisions on whether or how much to send their children to school. Note that some school inputs may have no direct effect on child attainment but will still be productive if they make children attend more regularly. Other potential school inputs will not be used if they are deemed unproductive, particularly if they lower school attendance. This is true even if the input is available at no cost to the school, as would be the case if the school receives a fixed set of materials from a central authority. The inputs actually used in practice will reflect the decisions of the local school administrator.[62]

Even this simple description of the human capital production process illustrates the complications from interpreting an empirical approximation to Eq. (6.1). Data sets will typically not include information on student ability, and so μ_k will be part of the error term. Because school inputs will be selected to complement student abilities, $\mathrm{Cov}(\mu_k, Z_k) \neq 0$ which will bias the coefficients on school inputs. As an example, suppose that schools that assign more independent projects do so only if their students have superior levels of ability, μ_k. The estimated effect of independent study on cognitive attainment will be positive, but one cannot infer that all schools should assign independent work.

Parents will send their children to the school more or less intensively depending on the mix of school inputs. Information on student cognitive attainment will only be available for children who attend on the day of the test, and they will presumably be the children whose abilities are most complementary with the mix of school inputs used. This nonrandom sorting into the school will further bias the estimates. Suppose, for example, that better schools retain a high fraction of their students while weaker schools retain only students from the upper tail of the ability distribution. The distribution of average test scores across schools will understate the true productivity of the superior schools.

Measures of the school inputs Z_k will typically include budgeted personnel, materials, and equipment, but the actual utilization of these inputs will reflect the decisions of the teachers and the school administration. Textbooks may be assigned but not used. Teaching methods may be prescribed but not followed. Teachers may be paid but not present. Severe measurement errors will exist because of the deviation between inputs allocated and used.

[62] The prices could be viewed as shadow prices of time if inputs require the administrator to provide training in their application. For example, the price of applying a new teaching method would include the cost of the in-service training as well as any associated materials.

Equation (6.2) suggests that schools will use inputs more intensively if they complement household assets or child abilities. Some combinations of inputs may be used in wealthier communities and others in poorer communities; other combinations with high ability children and still others with low ability populations. Natural complementarities among inputs imply that they will not be selected independent of one another but are selected in combination. If inputs are selected in groups and not independently, the frequent practice of using a very large number of highly intercorrelated school attributes in estimating educational production functions will not reliably measure the independent productivity of the included inputs.

Our model of input choice presumes that the local school manager wishes to maximize students' cognitive attainment. This may well be the case in schools with considerable parent involvement so that the school manager is acting as an agent for the parents. However, the principal may not serve as an agent for the parents. Politically or economically prominent local individuals may influence the school to misapply inputs to meet other ends, resulting in resources being diverted to unneeded school construction or hiring relatives. In other settings, the objective may be to maximize the utility of the teachers, resulting in frequent absenteeism. If the school administrator's objective function differs across schools, the observed pattern of inputs and cognitive attainment cannot be described by a single production function that assumes inputs are being allocated to maximize cognitive attainment subject to the mix of inputs.

Use of longitudinal data on test scores has been proposed as a possible correction for unobserved ability. If the growth in cognitive achievement is proportional to the initial test score, a regression of the change in test scores on the initial test score plus measures of A_i and Z_j may correct for the unobserved ability. However, μ_k would be expected to affect child time in school during the year, and so unobserved ability will still affect how intensively school inputs are used during the school year. Consequently, differencing will not purge the regression of ability bias. Furthermore, test scores measure cognitive ability with error, so taking the difference between two test scores may result in a high ratio of noise to the actual change in cognitive ability, and the use of the first test score as a regressor creates measurement error bias.[63]

Household surveys provide information on siblings and other members of the household. Various studies have used this feature of household data to generate household fixed effects to control for unobserved health or ability endowments and other unobserved household variables.[64] Data on twins, for example, are thought to be ideal for

[63] For more detail on these and other problems related to estimation of education production functions, see the careful review by Todd and Wolpin (2003).

[64] Household surveys usually collect detailed information only on the children residing in the household, ignoring the schooling of children who have been fostered out or who have left the household. The resulting selection bias can be large, but in an uncertain direction. Children who stay in school longest may have to leave the household for secondary schooling in larger towns. Alternatively, children who quit school and enter the labor market at early ages may leave the household first. The decision to leave the household might also be conditoned on unobserved abilities, but it is unclear whether the most able would leave at younger or older ages.

controlling for differences in natural ability (e.g., Behrman and Rosenzweig, 2002). Absent twins data, other studies include data on all children per household but estimate their models with household fixed effects (Lillard and Willis, 1994; Glewwe, Jacoby and King, 2001, and Section 4.2 of Strauss and Thomas in this Handbook). Such methods have the potential of correcting for ability bias on cognitive development, but they still leave an uncomfortable presumption that remaining differences in schooling intensity across children are not a result of child-specific attributes (Todd and Wolpin, 2003).

Randomized input assignments of the type reviewed by Duflo, Glennerster and Kremer in this Handbook may be a mechanism to avoid the endogenous school input problem. Of course, Eq. (6.2) suggests that the local school authority will decide whether or how the inputs are used, and parents may alter the child's time in school as well, and so there will be variation in input utilization across schools receiving the inputs. In addition, the productivity of the input (say, a new English language textbook) will depend on the other inputs available to the school (e.g., trained teachers, English medium instruction). An experimental infusion of certain inputs, but not others, may succeed in some settings, complicating the applicability of the lessons to other schools and contexts. As an example, Glewwe, Kremer and Moulin (2003) find that making textbooks more available in Kenya benefited students in the upper tail of the ability distribution who were prepared for the English medium texts, but the texts had no impact on below-average students who could not understand English. It is plausible that the textbooks encouraged more able students to attend school more regularly while discouraging weaker students from attending.

Studies that compare the productivity of school types, such as public versus private schools or religious versus secular schools, face similar challenges. If parents can choose where to send their children to school as well as how intensively to invest in schooling, then parents will select the school inputs Z_k and $\bar{Z}$ jointly with the type of school. These inputs will be correlated with all of the households assets and child ability endowments, and so differences in average school outcomes across school types will reflect the nonrandom sorting of children across school types. It is possible to measure how the availability of new choices affects schooling outcomes as in the Angrist et al. studies discussed in Section 3.4. However, this is the joint effect of attending the school and the application of home and school inputs, and not a means of establishing how individual inputs affect student outcomes.

No study has been able to address all of the difficulties outlined above, and problems of selection, measurement error, and differential objective functions are likely to be larger in developing than developed country settings. The many studies of the educational production process in developing countries have failed to generate consistent findings.[65] Teacher or school attributes that appear critically important for student performance in one study prove unimportant or even a detriment in another. This lack of

[65] For reviews of the results of educational production function estimation, see Hanushek (1997), Glewwe (2002), and Hanushek and Luque (2003).

consistency stands in marked contrast to the high degree of consistency across studies of estimated private returns to schooling which has been used to validate the use of log earnings equations in studies of returns to human capital. We conclude that the various misspecification problems that arise when estimating equations such as (6.1) are so pervasive as to render the coefficient estimates unreliable.

6.2.1. School quality

In many studies of schooling demand, one or more measures of the relative quality of local schools prove to be significant in explaining cross-sectional variation in attendance or enrollment, controlling for other household, school and community factors.[66] This contrasts with the lack of consensus regarding the impact of school quality on student achievement. This may suggest that the way school quality matters for learning is more in attracting and retaining children in school than in changing the pace of learning across children in different schools. Equally plausible is that the biases inherent in estimating the total derivative $\frac{\partial q}{\partial Z_{ik}}$ in (6.2) have led to the lack of consensus regarding the impact of school quality on student achievement, but that estimates of the impact of school quality on child time use ($\frac{\partial H_k}{\partial Z_{ik}}$ in (6.2)) are more reliable.

Before pressing this conclusion harder, it should be noted that the typical study of school quality on test scores uses numerous measures of school quality while the typical study of school quality on schooling demand uses only one or two measures of school quality. It may be that the perception that school quality is more important for child time in school than for cognitive development is due to more selective concentration on school quality measures that yield expected results in the schooling demand literature. Again, there is not a sufficiently well-developed literature on the relative importance of school quality on attendance versus school achievement to assess the reliability of the common finding that school quality affects child time in school. A particularly fruitful area for future research would be to evaluate whether variation in school quality matters more for explaining variation in attendance than for test scores, using a single data set appropriate for both questions.

6.2.2. Promotion, repetition and dropping out

In the absence of student test scores, some researchers and policymakers refer to repetition, promotion and dropping out as measures of student success or failure. This view is wrong because repetition and promotion rates reflect a confluence of factors, including academic standards and how strictly they are enforced, student performance and motivation, and household demand. Some schools and school systems practice social

[66] See for examples Case and Deaton (1999), Alderman, Orazem and Paterno (2001), Bedi and Marshall (2002), and Glick and Sahn (2006).

promotion in which all children are promoted to the next grade regardless of academic performance.[67] The belief underlying this practice is that children who are not promoted become discouraged and withdraw from school. The counter argument is that ill-prepared children fall further and further behind their peers and become discouraged, whereas holding them back until they are ready to progress allows time for remediation and eventual academic progress.

Obtaining measures of grade repetition and dropping out behavior is also difficult both at the individual and cohort levels. Of the three sample survey questionnaires we cited above, only one (the IFLS questionnaire) contains an attempt to determine the timing of school leaving so that it is possible to reconstruct the period when a child who is no longer enrolled was a student. By using the reported highest grade completed together with the child's current age, it is possible to get a rough estimate of age of entry in school; but, without direct questions about age of entry in school, number of grades repeated and which grades were repeated, whether a child who is currently in school had ever stopped for a year, and when a child who is no longer enrolled dropped out of school, it is not possible to characterize accurately the schooling histories of children in different settings and how these have changed across cohorts. Surveys such as the Malaysia Family Life Surveys, the Indonesia Family Life Surveys, and the Cebu Longitudinal Health and Nutrition Study (Philippines) have collected detailed life history data that contain these pieces of information. Even considering problems with the recall ability of the respondent and the availability of supporting documents which determine the reliability of the elicited information, such histories provide better data for understanding the schooling process.

6.3. Measurement of exogenous variables

Empirical tests of the theoretical static model of child schooling require data on child time allocations, household, school and community attributes, measures of the value of child time, and measures of schooling outcomes. This subsection discusses alternative empirical measures for the variables needed to estimate (4.8).[68]

6.3.1. Direct price of schooling

The direct costs faced by a household in sending a child to school include expenditures for tuition, required books and educational materials, transportation, uniform, exams,

[67] King, Orazem and Paterno (2002) examine whether social promotion increased the probability that a child continued in school in the Northwest Frontier Province of Pakistan. They find that promotions based on attendance and performance on tests increased the probability of continuation, but that promotions that were unrelated to school performance had virtually no effect on continuation. Consequently, continuation may be a better indicator of school performance than is promotion.

[68] Behrman and Rosenzweig (1994) discuss the implications of uneven data quality across developing countries for cross-country macroeconomic analyses of human capital and growth.

and admission fees required to gain access to the school. These costs usually vary by type of school. Private schools typically charge more than government schools, but there may be cost variation across private schools and across government schools as well. As noted in the previous subsection, variation in school distance (and thus in the cost of transportation) has proved to be a useful shifter of child time in school in several different settings, but distance to the nearest school for any one household or child is endogenous for at least two reasons – school choice which may be accompanied by a household relocating closer to a school, and endogenous school placement which is presumably aimed to reduce average distance to a school, so these observed marginal effects of school distance are likely biased downward.[69]

Market measures of the price of schools have typically included the average or median distance or cost across the schools that are available to the household (Gertler and Glewwe, 1990; Lavy, 1996; Alderman, Orazem and Paterno, 2001; Bedi and Marshall, 2002; Glick and Sahn, 2006). It is critically important that the price of schooling is not the price actually paid by the household but is a market-level measure of the variation in the price of schools (such as the average price in the village), since the actual price of the school selected is endogenous. This can lead to counterintuitive results if parents choose a more costly school that is a better school over nearer schools that are not as good. For example, Brown and Park (2002) found that increasing distance to school and the school price lowered the probability of dropout. An alternative approach suggested by Lavy (1996) is to estimate school price using as instruments area-specific characteristics that affect school placement but are not correlated with the demand for schooling. For Ghana, he uses distance to a public telephone and the post office to help identify the distance to a middle school in a given community; but assuming that community-specific attributes do not affect both the demand for schooling and its supply in the community is likely to be contested by some.

Average school price elasticities tend to be small, ranging from 0 to −0.4 in the studies mentioned above. They are of larger magnitude for private schools than for government schools. Poorer households are more responsive to price than richer households in all of the studies. The one study we found with comparable urban and rural estimates (Bedi and Marshall, 2002) did not find significant differences in price elasticities between urban and rural areas of Honduras, but it would be difficult to generalize from that study. More recently, several African countries have dispensed with school fees.

[69] Filmer (2004) provides a summary look at the availability of schools in several developing countries, as measured by the average distance to the nearest school reported in Demographic and Health Surveys. The data show large variation in average distance to school for households across the countries – from just 0.2 kilometers in Bangladesh to over seven kilometers in Chad – but the median average distance among 21 countries is only 1.4 km. The range in average distance to secondary schools is much greater – from 1.8 kilometers in Bangladesh to over 71 kilometers in Mali. Filmer (2004) finds that the distance to the nearest school is statistically significantly negatively related to school enrollment in 11 out of the 21 countries at the primary level (at the 5 percent significance level), and in six out of 21 countries at the secondary level. These results do not provide an overwhelming case for new school construction programs.

Deininger (2003) reports substantial increases in school enrollments in Uganda following the fee reduction. Ignoring the possibility that the fee reduction was anticipated ahead of its adoption, he nonetheless finds a substantial reduction in enrollment gaps between boys and girls, between urban and rural areas, and between rich and poor that suggest more price elastic demand for girls, for rural residents and for the poor.

6.3.2. *Opportunity cost of child time*

Most child labor occurs in the context of a family enterprise without pay, and so most of child labor occurs in the informal sector. This complicates the measurement of the opportunity cost of child time in school. One option (Orazem, 1987; Jacoby, 1993; Glewwe and Jacoby, 2004) is to estimate a production function and compute the marginal product of child labor. This method, however, has an added complexity in that the marginal product of time is a function of time spent in home production, and is consequently a function also of time spent in school. The endogeneity of the marginal product of time is a difficult problem to solve in that the general solution to the household's optimization problem makes child time in school a function of all the exogenous variables in the system. A tractable alternative is to approximate the child time equations (4.6) by assuming that factors such as household productive assets or technologies shift the child's marginal product of time. If, for example, we assume that child marginal productivity takes the form $Q_{H_i} = Q(A, \tau, \zeta)$ where ζ is a random error, then the reduced-form time use can be approximated by $H_i = H_i(P_C, P_f, P_m, A, Z, \tau, d_f, \mu_i)$ without having to impose the structural relationship among Q_{H_i}, A, and τ. For most applications, the reduced-form relationship will be sufficient.

An alternative is to use measures of local child market wages. However, in households where children work at home in the face of an active local child labor market, it must be true that the marginal product of child time in the home enterprise exceeds the market child wage, and so the average wage would be a lower bound measure of the value of child time. More problematic is that even where child labor markets are active, the majority of children do not work for a wage and so the observed average wage is subject to serious selection bias problems that may outweigh any benefit of using wages as a measure of child time.

It is plausible that the value of child time varies with age and physical maturity and not with what the child is learning in school, at least for the primary cycle. Several pieces of evidence suggest this conclusion. Rosenzweig (1980) finds that agricultural day wages in India did not rise with education. Fafchamps and Quisumbing (1999) observe that even adult returns to agricultural production in Pakistan rose with measures of the health and physical stature of both the husband and the wife in farm families, but not with their education. This suggests that age could be used as a proxy for the value of time, except that capacity for human capital accumulation will also rise with age, and so using age to distinguish between opportunity costs of schooling and capacity for learning can be problematic. Nevertheless, exogenous variation in age due to local policies concerning when a child can start or leave school or when a child can legally

initiate work can create exogenous variation in the opportunity costs of schooling across markets due to factors outside the household control.[70]

7. Conclusions

Compared to developed countries, developing countries have large gaps in schooling attainment between men and women and between urban and rural residents. As countries develop and education levels rise, these gaps have tended to close. Most notable has been the reduction in educational gaps between males and females in many countries. Nevertheless, large gaps within countries remain. Comparisons across countries show that such gaps can be substantially different even among countries at similar income levels.

Certainly, governments believe policies and interventions matter, as they invest considerable resources in providing or financing education for their citizens. The potential private returns of education to citizens and the potential social returns to the state require that these investments pay off in terms of more academic achievement and higher productivity. This is not always assured, in part because time spent in school may not lead to a commensurate increase in knowledge and skills. For this reason, governments, international agencies and individual researchers are devoting more attention to examining which policies and programs are most cost-effective. Establishing the impact of these policies and programs requires an understanding of the incentives and constraints faced by all parties involved, the school providers, the parents and the children.

Increasingly, assessments are making use of experimental methods to estimate the impact of policy changes and programs on educational attainment. Such methods provide counterfactual populations for treated groups. This chapter argues that in order to generalize from individual experimental evidence to other settings, it is important to frame the experimental outcomes in the context of a behavioral model that can help explain why a particular pilot project generates the outcomes that it does. In addition, behavioral models are needed to guide nonexperimental evaluations in contexts where experimental evaluations are infeasible.

Behavioral models can also help to frame the magnitude of experimental results. Knowing that pilots are likely to be placed in settings where they will be atypically successful suggests that experimental outcomes derived from such settings will overstate the outcomes for the population as a whole.[71] Understanding that income surprises will

[70] Angrist and Krueger (1991), Tyler (2003), and Gunnarsson et al. (2006) use variation in child labor and truancy laws to create a source of exogenous variation in child ages in their estimates of schooling demand.

[71] For example, Miguel and Kremer's (2004) analysis of deworming medicine discussed extensively in the chapters written by Miguel and Glewwe and by Mwabu in this Handbook found large impacts on child health and school attendance in areas where 92% of the children were infected. One would expect more modest results if the program were to be expanded to areas with more modest infection rates. See also the discussion of vouchers in Section 3.

have atypically large impacts on credit constrained households suggests that the income elasticities derived by measuring short-term schooling demand responses to transitory income shocks will exaggerate the schooling demand responses to permanent income changes.

There are several important avenues for future research:

(a) A common, albeit not universal finding, is that the elasticity of child schooling with respect to household income is larger in developing than in developed countries. According to the Becker–Tomes framework, credit-constrained households will concentrate their transfers of wealth to their children in the form of human capital, but the investments will be below the efficient levels. Public provision of schooling should reduce the degree of under-investment and help to decrease the intergenerational transmission of poverty. Future research could evaluate schooling demand in developing countries in light of the Becker–Tomes model: Is it true that the tie between parent's income and child schooling is indeed greater in developing countries, and does the current allocation of schooling reduce or perpetuate income inequality in developing countries?

(b) Studies that examine household reactions to income shocks have followed behaviors for a limited time after the shock. Studies also tend to look only at consumption responses or time allocation responses in isolation rather than examine them as simultaneous decisions as implied by Eqs. (5.6a) and (5.6b). Longer time-series data with more complete model specifications would help to disentangle the short- from long-run schooling effects and time allocation from consumption responses. The same plea for future research applies with respect to household reactions to permanent and transitory price shocks.

(c) Longitudinal analysis of cognitive attainment is needed to establish whether lost human capital from transitory increases in child labor or school absences due to adverse income shocks is reversible or permanent.

(d) Because many of the interesting decisions regarding schooling investments do not lend themselves easily to experimental methods, more studies are needed that use convincing instruments to correct for the likely endogeneity of child time or of measurement error in explanatory variables used in evaluating schooling outcomes.

(e) Despite the large differences in schooling attainment between urban and rural residents, there are very few studies that use similar methods to compare schooling demand across urban and rural areas. As rural to urban migration becomes more important, rural migrants will need skills more similar to those of their urban contemporaries. Returns to rural education should not be undervalued because we fail to include the consequences of rural schooling on rural-urban migration and the resulting increases in lifetime earnings that result.

(f) There are no studies that evaluate the impact of urban–rural schooling gaps on economic development, or whether there are important external benefits to raising schooling outcomes in rural areas. Improved schooling in rural areas has been shown to increase outmigration from rural to urban markets, to supply skilled labor needed to fuel industrialization, to increase the pace of technological adoption and yield gain by farmers, and to reduce earnings inequality between urban and rural markets. If, as we

suspect, these returns to schooling are large and of critical importance to a country's transition from an agrarian to an industrialized economy, we should find that improved rural schooling is tied to economic growth.

(g) Most of the studies reviewed in this chapter have focused on years of schooling and enrollment rates as measures of educational achievement. The economics research using measures of learning achievement, such as standardized test scores, has been on the rise, but this is generally still a nascent area. The growth of this body of research depends greatly on more countries having test score data that are linked to household and child characteristics and that are available for more than one year.

(h) As data and evidence from experimental evaluations in different settings become more available, there will be two important roles linking experiments to behavioral models:

1. Behavioral models can be tested using experimental methods to examine the extent to which econometric models can replicate the results from the experiments. Suitably calibrated behavioral models can then be used to conduct simulated counterfactual policy exercises as demonstrated by Todd and Wolpin (2006).
2. Behavioral models can be used to aggregate the results of various experiments to predict how similar interventions will perform in alternative environments.

(i) In many countries, girls receive less schooling than boys. In virtually all of those countries, least squares estimates of private returns to schooling are lower for boys than for girls. This could be a consequence of parents applying a larger discount to girls' future earnings. It could also reflect a higher opportunity cost of schooling for girls than boys. However, it may also be a consequence of greater selection bias in estimated returns to girls' schooling if the women who ultimately enter wage labor are disproportionately able compared to the average male wage worker. If girls' schooling also has greater external social benefits from reduced fertility or improved family health as many have argued, then there is a *prima facie* case that these developing countries under-invest in girls' education. However, except for Schultz (1988), economists have not conducted a systematic evaluation of the relative size of possible biases in estimated returns to schooling between males and females in developing countries. Countries such as Mexico have implemented conditional transfer programs that allocate slightly larger grants to girls than to boys on the presumption that girls' schooling responds less elastically to transfers and that returns to the investment are greater. The very limited evidence available does not refute that assumption, but much more work needs to be done.

(j) Average returns to schooling for rural nonmigrant residents are lower than for urban residents. On-farm returns to schooling are lower than off-farm returns to schooling in areas using traditional agricultural methods. However, it is a mistake to conclude that rural investments in schooling are unimportant. More educated rural residents migrate in greater numbers to urban markets, suggesting that observed rural returns to schooling are understated. We need studies that establish the returns to rural education for migrants versus nonmigrants.

Appendix Table 1
Country household data sets used to estimate returns to schooling in Figs. 4 and 5

Country	Years
Albania (ALB)	2002
Argentina (ARG)[a]	1994, 2001
Bangladesh (BNG)	2000
Bolivia (BOL)	2002
Brazil (BRA)	1995, 2001
Burkina Faso (BFA)	1994, 2003
Bulgaria (BUL)	2001
Burundi (BUR)	1998
Cambodia (KHM)	1997, 2004
Chile (CHL)	1990, 1996
Colombia (COL)	1995, 2000
Costa Rica (CRI)	1995, 2000
Croatia (HRV)	2004
Dominican Republic (DOM)	1997, 2004
East Timor (TMP)	20 01
Ecuador (ECU)	1995, 2004
Egypt (EGY)	1998
El Salvador (SLV)	1 995, 2002
Estonia (EST)	2000
Ghana (GHA)	1991
Guatemala (GUA)	1989, 2002
Guyana (GUY)	1992
Haiti (HTI)	2001
Honduras (HON)	1995, 2003
Indonesia (IDN)	2002
Jamaica (JAM)	1996, 2002
Jordan (JOR)	2002
Kyrgyzstan (KGZ)	1996
Maldives (MDV)	1998
Mauritania (MRT)	2000
Mexico (MEX)	1994, 2002
Mongolia (MON)	2002
Morocco (MAR)	1991
Mozambique (MOZ)	1996
Nicaragua (NIC)	1993, 2001
Nigeria (NGA)	2003
Pakistan (PAK)	1999, 2001
Palau (PLW)[a]	2000
Panama (PAN)	1995, 2003
Paraguaqy (PGY)	1995, 2001
Peru (PER)	1994, 2002
Romania (ROM)	1994
Sao Tome and Principe (STP)	2000
Sierra Leone (SLE)	2003
Tajikistan (TJK)	2003
Thailand (THA)	1990, 1994, 2002

(*continued on next page*)

Appendix Table 1
(*continued*)

Country	Years
Uganda (UGA)	2002
Uruguay (URY)[a]	1995, 2003
Vietnam (VNM)	1992

Source: Fares, Montenegro and Orazem (2007). Country three-letter codes used in Figs. 2 and 3 are reported in parentheses.
[a]Data not available for separate estimates for urban and rural areas.

References

Abdulai, A., Huffman, W.E. (2005). "The diffusion of new agricultural technologies: The case of crossbreeding technology in Tanzania". American Journal of Agricultural Economics 87, 645–659.

Agesa, R.U. (2001). "Migration and the urban to rural earnings difference: A sample selection approach". Economic Development and Cultural Change 49 (4), 847–865.

Alderman, H., Kim, J., Orazem, P.F. (2003). "Design, evaluation, and sustainability of private schools for the poor: The Pakistan urban and rural fellowship school experiments". Economics of Education Review 22, 265–274.

Alderman, H., Orazem, P.F., Paterno, E.M. (2001). "School quality, school cost and the public/private school choices of low-income households in Pakistan". Journal of Human Resources 36, 304–326.

Alderman, H., Behrman, J.R., Lavy, V., Menon, R. (2001). "Child health and school enrollment: A longitudinal analysis". Journal of Human Resources 36 (1), 185–205.

Altonji, J.G., Dunn, T.A. (1996). "Using siblings to estimate the effect of school quality on wages". Review of Economics and Statistics 78 (4), 665–671.

Anderson, K.H., King, E.M., Wang, Y. (2002). "Market wages, transfers, and demand for schooling in Malaysia, 1976–1989". Journal of Development Studies 39 (3), 1–28.

Angrist, J.D., Bettinger, E., Kremer, M. (2006). "Long-term educational consequences of secondary school vouchers: Evidence from administrative records in Colombia". American Economic Review 96, 847–862.

Angrist, J.D., Krueger, A.B. (1991). "Does compulsory school attendance affect schooling and earnings?" Quarterly Journal of Economics 106, 979–1014.

Angrist, J.D., Krueger, A.B. (1999). "Empirical strategies in labor economics". In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3A. Elsevier Science B.V., Amsterdam.

Angrist, J.D., Bettinger, E., Bloom, E., King, E., Kremer, M. (2002). "Vouchers for private schooling in Colombia: Evidence from a randomized natural experiment". American Economic Review 92 (5), 1535–1559.

Anh, T.S., Knodel, J.E., Lam, D., Friedman, J. (1998). "Family size and children's education in Vietnam". Demography 35 (1), 57–70.

Antonovics, K.L., Goldberger, A.S. (2005). "Does increasing women's schooling raise the schooling of the next generation? Comment". American Economic Review 95, 1738–1744.

Becker, G.S. (1971). The Economics of Discrimination, second ed. The Univ. of Chicago Press, Chicago.

Becker, G.S. (1985). "Human capital, effort, and the sexual division of labor". Journal of Labor Economics 3 (1–2), S33–S58.

Becker, G.S. (1993). Human Capital, third ed. The Univ. of Chicago Press, Chicago.

Becker, G.S., Tomes, N. (1986). "Human capital and the rise and fall of families". Journal of Labor Economics 4, S12–S37.

Bedi, A.S., Edwards, J.H.Y. (2002). "The impact of school quality on earnings and educational returns–evidence from a low-income country". Journal of Development Economics 68 (1), 157–185.

Bedi, A.S., Gaston, N. (1999). "Using variation in schooling availability to estimate educational returns for Honduras". Economics of Education Review 18, 107–116.

Bedi, A.S., Marshall, J.H. (2002). "Primary school attendance in Honduras". Journal of Development Economics 69 (1), 129–153.

Behrman, J.R. (1996). Human Resources in Latin America and the Caribbean. The Johns Hopkins Univ. Press for the Inter-American Development Bank, Baltimore.

Behrman, J.R., Cheng, Y., Todd, P. (2004). "Evaluating pre-school programs when length of exposure to the program varies: A nonparametric approach". Review of Economics and Statistics 86 (1), 108–132.

Behrman, J.R., Rosenzweig, M.R. (1994). "Caveat emptor: Cross-country data on education and the labor force". Journal of Development Economics 44, 147–171.

Behrman, J.R., Rosenzweig, M.R. (2002). "Does increasing women's schooling raise the schooling of the next generation?" American Economic Review 92, 323–334.

Behrman, J.R., Rosenzweig, M.R. (2005). "Does increasing women's schooling raise the schooling of the next generation? Reply". American Economic Review 95, 1745–1751.

Behrman, J.R., Khan, S., Ross, D., Sabot, R. (1997). "School quality and cognitive achievement production: A case study for rural Pakistan". Economics of Education Review 16 (2), 127–142.

Besley, T., Case, A. (1993). "Modeling technology adoption in developing countries". American Economic Review 83, 396–402.

Bhalotra, S., Heady, C. (2003). "Child farm labor: The wealth paradox". World Bank Economic Review 17, 197–227.

Blundell, R., MaCurdy, T.E. (1999). "Labor supply: A review of alternative approaches". In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3. Elsevier Science B.V., Amsterdam.

Bommier, A., Lambert, S. (2000). "Education demand and age at school enrollment in Tanzania". Journal of Human Resources 35, 177–203.

Boucher, S., Stark, O., Taylor, J.E. (2007). "A gain with a drain? Evidence from rural Mexico on the New Economics of the brain drain". In press.

Brown, P.H., Park, A. (2002). "Education and poverty in rural China". Economics of Education Review 21 (6), 523–541.

Card, D. (1999). "The causal effect of education on earnings". In: Ashenfelter, O., Card, D. (Eds.), In: Handbook of Labor Economics, vol. 3A. Elsevier Science B.V., Amsterdam.

Card, D., Krueger, A.B. (1992). "Does school quality matter? Returns to education and the characteristics of public schools in the United States". Journal of Political Economy 100, 1–40.

Carneiro, P., Heckman, J. (2003). "Human capital policy". In: Heckman, J.J., Krueger, A.J. (Eds.), Inequality in America: What Role for Human Capital Policies?. MIT Press, Cambridge, MA.

Case, A., Deaton, A. (1999). "School inputs and educational outcomes in South Africa". Quarterly Journal of Economics 114 (3), 1047–1084.

Case, A., Paxson, C.H., Ableidinger, J. (2004). "Orphans in Africa: Parental death, poverty, and school enrollment". Demography 41 (3), 483–508.

Case, A., Yogo, M. (1999). "Does school quality matter? Returns to education and the characteristics of schools in South Africa". Working paper 7399. National Bureau of Economic Research.

Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., Rogers, F.H. (2006). "Missing in action: Teacher and health worker absence in developing countries". Journal of Economic Perspectives 20, 91–116.

Connelly, R., Zheng, Z. (2003). "Determinants of school enrollment and completion of 10 to 18 year-olds in China". Economics of Education Review 22 (4), 379–388.

Contreras, D. (2002). "Vouchers, school choice and the access to higher education". Discussion paper No. 845. Economic Growth Center, Yale University, June. Processed.

Cox-Edwards, A., Ureta, M. (2003). "International migration, remittances, and schooling: Evidence from El Salvador". Journal of Development Economics 72, 429–461.

Currie, J. (2001). "Early childhood education programs". Journal of Economic Perspectives 15 (2), 213–238.

Currie, J., Thomas, D. (2000). "School quality and the longer-term effects of head start". The Journal of Human Resources 35 (4), 755–774.

de Brauw, A., Rozelle, S., Zhang, L. (2005). "Labor market emergence and returns to education in rural China". Review of Agricultural Economics 27 (September), 418–424.

Deininger, K. (2003). "Does cost of schooling affect enrollment by the poor? Universal primary education in Uganda". Economics of Education Review 22, 291–305.

Duflo, E. (2001). "Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment". American Economic Review 91, 795–813.

Duflo, E. (2004). "The medium run effects of educational expansion: Evidence from a large school construction program in Indonesia". Journal of Development Economics 74 (1), 163–197.

Dunn, C.S. (2004). "Intergenerational transmission of lifetime earnings: New evidence from Brazil". University of Michigan, PhD dissertation.

Duraisamy, P. (2002). "Changes in returns to education in India, 1983–94, by gender, age cohort and location". Economics of Education Review 21 (6), 609–622.

Duryea, S., Hoek, J., Lam, D., Levison, D. (2006). "Dynamics of child labor: Labor force entry and exit in urban Brazil". In: Orazem, P.F., Sedlacek, G., Zafiris Tzannatos, P. (Eds.), Child Labor and Education in Latin America. Inter-American Development Bank, Washington, DC. In press.

Edmonds, E.V., Mammen, K., Miller, D.L. (2005). "Rearranging the family? Household composition responses to large pension receipts". Journal of Human Resources 40, 186–207.

Emerson, P.M., Souza, A.P. (2006). "Is child labor harmful? The impact of starting to work as a child on adult earnings". University of Colorado at Denver. Processed.

Fafchamps, M., Quisumbing, A. (1999). "Human capital, productivity, and labor allocation in rural Pakistan". Journal of Human Resources 34, 369–406.

Fares, J., Montenegro, C.E., Orazem, P.F. (2007). Variations in the Returns to Schooling Across and Within Developing Economies. The World Bank. Processed.

Filmer, D. (2004). "If you build it, will they come? School availability and school enrollment in 21 poor countries". Working paper 3340. World Bank Policy Research, June.

Folbre, N. (2006). "Measuring care: Gender, empowerment, and the care economy". Journal of Human Development 7 (2), 183–199.

Foster, A.D., Rosenzweig, M.R. (1996). "Technical change and human-capital returns and investments: Evidence from the green revolution". American Economic Review 86 (4), 931–953.

Foster, A.D., Rosenzweig, M.R. (2004). "Agricultural productivity growth, rural economic diversity, and economic reforms: India, 1970–2000". Economic Development and Cultural Change 52, 509–542.

Frankenberg, E., Smith, J.P., Thomas, D. (2003). "Economic shocks, wealth and welfare". Journal of Human Resources 38 (2), 280–321.

Friedman, M. (1957). A Theory of the Consumption Function. Princeton Univ. Press, Princeton, NJ.

Funkhouser, E. (1999). "Cyclical economic conditions and school attendance in Costa Rica". Economics of Education Review 18, 31–50.

Gertler, P., Glewwe, P. (1990). "The willingness to pay for education in developing countries: Evidence from rural Peru". Journal of Public Economics 42, 251–275.

Glewwe, P. (2002). "Schools and skills in developing countries: Education policies and socioeconomic outcomes". Journal of Economic Literature 40 (2), 436–483.

Glewwe, P., Hall, G. (1998). "Are some groups more vulnerable to macroeconomic shocks than others? Hypothesis tests based on panel data from Peru". Journal of Development Economics 56, 181–206.

Glewwe, P., Jacoby, H. (1995). "An economic analysis of delayed primary school enrollment in a low income country: The role of early childhood nutrition". Review of Economics and Statistics 77, 156–169.

Glewwe, P., Jacoby, H. (2004). "Economic growth and the demand for education: Is there a wealth effect?" Journal of Development Economics 74 (1), 33–51.

Glewwe, P., Jacoby, H., King, E. (2001). "Early childhood nutrition and academic achievement: A longitudinal analysis". Journal of Public Economics 81 (3), 345–368.

Glewwe, P., King, E.M. (2001). "The impact of early childhood nutritional status on cognitive development: Does the timing of malnutrition matter?". World Bank Economic Review 15 (1), 81–113.

Glewwe, P., Kremer, M., Moulin, S. (2003). "Textbooks and test scores: Evidence from a randomized evaluation in Kenya". Development Research Group. World Bank, Washington, DC.

Glewwe, P., Kremer, M., Moulin, S., Zitzewitz, E. (2004). "Retrospective vs. prospective analyses of school inputs: The case of flip charts in Kenya". Journal of Development Economics 74, 251–268.

Glick, P., Sahn, D.E. (2000). "Schooling of girls and boys in a West African country: The effects of parental education, income, and household structure". Economics of Education Review 19 (1), 63–87.

Glick, P., Sahn, D.E. (2006). "The demand for primary schooling in Madagascar: Price, quality, and the choice between public and private providers". Journal of Development Economics 79 (1), 118–145.

Godoy, R., Karlan, D.S., Rabindran, S., Huanca, T. (2005). "Do modern forms of human capital matter in primitive economies? Comparative evidence from Bolivia". Economics of Education Review 24, 45–53.

Grantham-McGregor, S.M., Walker, S.P., Chang, S.M., Powell, C.A. (1997). "Effects of early childhood supplementation with and without stimulation on later development in stunted Jamaican children". American Journal of Clinical Nutrition 66, 247–253.

Gunnarsson, V., Orazem, P.F., Sanchez, M.A. (2006). "Child labor and school achievement in Latin America". World Bank Economic Review 20, 31–54.

Handa, S. (2002). "Raising primary school enrollment in developing countries: The relative importance of supply and demand". Journal of Development Economics 69, 103–128.

Hanushek, E.A. (1997). "Assessing the effects of school resources on student performance: An update". Education Evaluation and Policy Analysis 19, 141–164.

Hanushek, E.A., Luque, J.A. (2003). "Efficiency and equity in schools around the world". Economics of Education Review 22, 481–502.

Haveman, R., Wolfe, B. (1995). "The determinants of children's attainments: A review of methods and findings". Journal of Economic Literature 23, 1829–1878.

Heckman, J.J., La Londe, R.J., Smith, J.A. (1999). "The economics and econometrics of active labor market programs". In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3. Elsevier, Amsterdam.

Heckman, J.J., Rubinstein, Y. (2001). "The importance of noncognitive skills: Lessons from the GED testing program". American Economic Review 91 (2), 145–149.

Hersch, J., Stratton, L.S. (1997). "Housework, fixed effects, and wages of married workers". Journal of Human Resources 32 (2), 285–307.

Holmes, J. (2003). "Measuring the determinants of school completion in Pakistan: Analysis of censoring and selection bias". Economics of Education Review 22, 249–264.

Hoxby, C.M. (2003). "School choice and school competition: Evidence from the United States". Swedish Economic Policy Review 10, 13–67.

Hsieh, C., Urquiola, M. (2006). "The effects of generalized school choice on achievement and stratification: Evidence from Chile's school voucher program". Journal of Public Economics 90, 1477–1503.

Huffman, W.E. (1977). "Allocative efficiency: The role of human capital". Quarterly Journal of Economics 91, 59–79.

Huffman, W.E., Orazem, P.F. (2006). "Agriculture and human capital in economic growth: Farmers, schooling and health". In: Evenson, R.E., Schultz, T.P., Pingali, P. (Eds.), Handbook of Agricultural Economics, vol. 3. North-Holland, Amsterdam.

Jacoby, H.G. (1993). "Shadow wages and peasant family labor supply: An econometric application to the Peruvian Sierra". Review of Economic Studies 60, 903–921.

Jacoby, H.G., Skoufias, E. (1997). "Risk, financial markets, and human capital in a developing country". Review of Economic Studies 64, 311–335.

Jacoby, H.G., Skoufias, E. (1998). "Testing theories of consumption behavior using information on aggregate shocks: Income seasonality and rainfall in rural India". American Journal of Agricultural Economics 80, 1–14.

Jalan, J., Glinskaya, E. (2003). "Improving primary school education in India: An impact assessment of DPEP-phase I". Indian Statistical Institute (New Delhi) and The World Bank. Processed.

James, E. (1993). "Why do different countries choose a different public-private mix of educational services?" Journal of Human Resources 28, 571–592.

Jimenez, E., Sawada, Y. (2001). "Public for private: The relationship between public and private school enrollment in the Philippines". Economics of Education Review 20, 389–399.

Jolliffe, D. (1998). "Skills, schooling, and household income in Ghana". World Bank Economic Review 12 (1), 81–104.

Joshi, S., Schultz, T.P. (2005). "Family planning as an investment in female human capital: Evaluating the long term consequences in Matlab, Bangladesh". Yale University. Processed.

Karoly, L.A., Greenwood, P.W., Everingham, S.S., Hoube, J., Kilburn, M.R, Rydell, C.P., Sanders, M., Chiesa, J. (1998). Investing in Our Children: What We Know and Don't Know About the Costs and Benefits of Early Childhood Interventions. RAND Corporation, Santa Monica, CA.

Kim, J., Alderman, H., Orazem, P.F. (1999). "Can private school subsidies increase enrollment for the poor? The Quetta urban fellowship program". World Bank Economic Review 13, 443–465.

King, E.M., Lillard, L.A. (1987). "Education policy and schooling attainment in Malaysia and the Philippines". Economics of Education Review 6 (2), 167–181.

King, E.M., Mason, A.D. (2001). Engendering Development through Gender Equality in Rights, Resources and Voice. Oxford Univ. Press, New York.

King, E.M., Orazem, P.F., Paterno, E.M. (2002). "Promotion with and without learning: Effects on student dropout". Iowa State University. Processed.

King, E.M., Orazem, P.F., Wohlgemuth, D. (1999). "Central mandates and local incentives: The Colombia education voucher program". World Bank Economic Review 13, 467–491.

Kochar, A. (2004). "Urban influences on rural schooling in India". Journal of Development Economics 74, 113–136.

Krueger, A.B., Lindahl, M. (2001). "Education for growth: Why and for whom?". Journal of Economic Literature 39, 1101–1136.

Langemeier, M.R., Patrick, G.F. (1990). "Farmers' marginal propensity to consume: An application to Illinois grain farms". American Journal of Agricultural Economics 72, 309–316.

Langemeier, M.R., Patrick, G.F. (1993). "Farm consumption and liquidity constraints". American Journal of Agricultural Economics 75, 479–484.

Lavy, V. (1996). "School supply constraints and children's educational outcomes in rural Ghana". Journal of Development Economics 51, 291–314.

Lillard, L.A., King, E.M. (1984). "Methods for analyzing schooling choice with household survey data". Report N-1963-AID. The RAND Corporation, Santa Monica, CA.

Lillard, L.A., Willis, R.J. (1994). "Intergenerational educational mobility: Effects of family and state in Malaysia". Journal of Human Resources 29 (4), 1126–1166.

Lozoff, B., Jimenez, E., Hagen, J., Mollen, E., Wolf, A.W. (2000). "Treatment for iron deficiency in infancy poorer behavioral and developmental outcome more than 10 years after". Pediatrics 105 (4), 51–62.

MaCurdy, T.E. (1985). "Interpreting empirical models of labor supply in an intertemporal framework with uncertainty". In: Heckman, J., Singer, B. (Eds.), Longitudinal Analysis of Labor Market Data. Cambridge Univ. Press, Cambridge.

Maliki, O. (2005). "Education policy and intergenerational transfers in Indonesia". PhD dissertation. University of Hawaii, Unpublished.

Martinelli, C., Parker, S.W. (2003). "Should transfers to poor families be conditional on school attendance". International Economic Review 44 (2), 523–544.

Martorell, R. (1999). "The nature of child malnutrition and its long-term implications". Food and Nutrition Bulletin 20 (3), 288–292.

McEwan, P.J. (2001). "The effectiveness of public, catholic, and non-religious private schools in Chile's voucher system". Education Economics 9 (2), 103–128.

Miguel, E., Kremer, M. (2004). "Worms: Identifying impacts on education and health in the presence of treatment externalities". Econometrica 72 (1), 159–217.

Mincer, J. (1974). Schooling, Experience and Earnings. Columbia Univ. Press, New York.

Mizala, A., Romaguera, P. (2000). "School performance and choice: The Chilean experience". Journal of Human Resources 35 (2), 392–417.

Moffitt, R. (1996). "Symposium on school quality and educational outcomes". Review of Economics and Statistics 78, 559–561.

Munshi, K., Rosenzweig, M. (2006). "Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy". American Economic Review 96 (4), 1225–1252.

Neri, M.C., Gustafsson-Wright, E., Sedlacek, G., Orazem, P.F. (2006). "The responses of child labor, school enrollment, and grade repetition to the loss of parental earnings in Brazil, 1982–1999". In: Orazem, P.F., Sedlacek, G., Zafiris Tzannatos, P. (Eds.), Child Labor and Education in Latin America. Inter-American Development Bank, Washington, DC. In press.

Orazem, P.F. (1987). "Black–white difference in schooling investment and human capital production in segregated schools". American Economic Review 77, 714–723.

Parente, S.L., Rogerson, R., Wright, R. (2000). "Homework in development economics: Household production and the wealth of nations". Journal of Political Economy 108, 680–687.

Patrinos, H.A., Psacharopoulos, G. (1997). "Family size, schooling and child labor in Peru – An empirical analysis". Journal of Population Economics 10, 387–405.

Paxson, C.H. (1992). "Using weather variability to estimate the response of savings to transitory income in Thailand". American Economic Review 82, 15–33.

Paxson, C.H. (1993). "Consumption and income seasonality in Thailand". Journal of Political Economy 99, 39–72.

Plug, E., Vijverberg, W. (2003). "Schooling, family background, and adoption: Is it nature or is it nurture?" Journal of Political Economy 111, 611–641.

Pollitt, E. (2000). "Developmental sequel from early nutritional deficiencies: Conclusive and probability judgments". Journal of Nutrition 130 (2), 350S–353S.

Quisumbing, A.R., Maluccio, J.A. (2000). "Intrahousehold allocation and gender relations: New empirical evidence from four developing countries". Discussion paper No. 84. International Food Policy Research Institute, FCND (April).

Ravallion, M., Wodon, Q. (2000). "Does child labor displace schooling? Evidence on behavioral responses to an enrollment subsidy". Economic Journal 110, 158–175.

Reid, M.G. (1952). "Effect of income concept upon expenditure curves of farm families". In: Studies in Income and Wealth, vol. 15. National Bureau of Economic Research, New York.

Rosati, F.C., Rossi, M. (2003). "Children's working hours and school enrollment: Evidence from Pakistan and Nicaragua". World Bank Economic Review 17, 283–295.

Rosen, S. (1977). "Human capital: A survey of empirical research". In: Ehrenberg, R.G. (Ed.), Research in Labor Economics, vol. I. JAI Press, Greenwich, CT.

Rosenzweig, M.R. (1980). "Neoclassical theory and the optimizing peasant: An econometric analysis of market family labor supply in a developing country". Quarterly Journal of Economics 94, 31–55.

Rosenzweig, M.R., Wolpin, K.I. (1986). "Evaluating the effects of optimally distributed public programs: Child health and family planning interventions". American Economic Review 76, 470–482.

Rosenzweig, M.R., Wolpin, K.I. (1988). "Migration selectivity and the effects of public programs". Journal of Public Economics 37, 265–289.

Schulz, T.P. (1988). "Education investments and returns". In: Chenery, H., Srinivasan, T.N. (Eds.), Handbook of Development Economics, vol. 1. North-Holland, Amsterdam.

Schultz, T.P. (2001). "Women's roles in the agricultural household: Bargaining and human capital investments". In: Gardner, B.L., Rausser, G.C. (Eds.), Handbook of Agricultural Economics. North-Holland, Amsterdam.

Schultz, T.P. (2002). "Why governments should invest more to educate girls". World Development 30, 207–225.

Schultz, T.P. (2004). "School subsidies for the poor: Evaluating the Mexican Progresa poverty program". Journal of Development Economics 74, 199–250.

Schultz, T.W. (1964). Transforming Traditional Agriculture. Yale Univ. Press, New Haven, CT.

Schultz, T.W. (1975). "The value of the ability to deal with disequilibria". Journal of Economic Literature 13, 827–846.

Singh, I., Squire, L., Strauss, J. (Eds.) (1986). Agricultural household models. The Johns Hopkins Univ. Press, Baltimore, MD.

Smeeding, T., Weinberg, D. (2001). "Toward a uniform definition of household income". Review of Income and Wealth 47, 1–24.

Solon, G.R. (2002). "Cross-country differences in intergenerational earnings mobility". Journal of Economic Perspectives 16, 59–66.

Strauss, J., Thomas, D. (1995). "Human resources: Empirical modeling of household and family decisions". In: Srinivasan, T.N., Behrman, J.R. (Eds.), Handbook of Development Economics, vol. 3A. North-Holland, Amsterdam, pp. 1883–2023.

Strauss, J., Thomas, D. (1998). "Health, nutrition, and economic development". Journal of Economic Literature 36 (June), 436–482.

Tansel, A. (1997). "Schooling attainment, parental education, and gender in Cote d'Ivoire and Ghana". Economic Development and Cultural Change 45 (4), 825–856.

Tansel, A. (2002). "Determinants of school attainment of boys and girls in Turkey: Individual, household and community factors". Economics of Education Review 21, 455–470.

Thomas, D., Strauss, J. (1992). "Prices, infrastructure, household characteristics and child height". Journal of Development Economics 39 (2), 301–331.

Thomas, D., Beegle, K., Frankenberg, E., Sikoki, B., Strauss, J., Teruel, G. (2004). "Education in a crisis". Journal of Development Economics 74, 53–85.

Todd, P., Wolpin, K.I. (2003). "On the specification and estimation of the production function for cognitive achievement". Economic Journal 113 (1), F3–F33.

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

Tyler, J.H. (2003). "Using state child labor laws to identify the effect of school-to-work on high school achievement". Journal of Labor Economics 21, 381–408.

Welch, F. (1970). "Education in production". Journal of Political Economy 78, 35–59.

Wolpin, K.I. (1982). "A new test of the permanent income hypothesis: The impact of weather on the income and consumption of farm households in India". International Economic Review 23, 583–594.

Wooldridge, J.M. (2002). Econometric Analyses of Cross Section and Panel Data. The MIT Press, Cambridge, MA.

Yang, D. (1997). "Education and off-farm work". Economic Development and Cultural Change 45, 613–632.

Yap, Y.-T., Orazem, P.F., Sedlacek, G. (2006). "Limiting child labor through behavior-based income transfers: An experimental evaluation of the PETI program in rural Brazil". In: Orazem, P.F., Sedlacek, G., Zafiris Tzannatos, P. (Eds.), Child Labor and Education in Latin America. Inter-American Development Bank, Washington, DC. In press.

Zimmerman, F.J. (2003). "Cinderella goes to school". Journal of Human Resources 38, 557–591.