

Expanding Access and Increasing Student Learning in Post-Primary Education in Developing Countries: A Review of the Evidence

Abhijit Banerjee, Paul Glewwe, Shawn Powers, and Melanie Wasserman¹

Abdul Latif Jameel Poverty Action Lab (J-PAL)
Post-Primary Education Initiative Review Paper²

Version of 2013-04-09 – Comments Welcome

I. Introduction

Effective, evidence-based policies on post-primary education are of vital importance as many developing countries start to see a bulge in secondary and post-secondary enrollment, the product of the achievement of near-universal access to primary school. Finding ways to deliver and promote access to high-quality post-primary education, and to ensure that education is relevant to labor market needs, is one of the great challenges of our times. This must be accomplished in countries where governments face severe budget constraints and many, of not most, parents are too poor to cover the costs out of pocket. International reports such as *A Global Compact on Learning*, by the Center for Universal Education at the Brookings Institution, emphasize providing opportunities for post-primary education as a first-tier policy challenge (Center for Universal Education, 2011). In addition, there has been considerably less progress in gender parity at the secondary level. Meeting these challenges will require a combination of using existing resources more effectively – which requires both understanding which inputs are key and which are not—and a range of innovations that may fundamentally alter the current methods of instruction.

To that end, the Abdul Latif Jameel Poverty Action Lab (J-PAL) has launched a Post-Primary Education Initiative intended to promote policy-relevant research on secondary and post-secondary education in developing countries, which together will be referred to as post-primary education. This paper is a first step in that process. We review the evidence to date on post-primary education and highlight the gaps in the literature, with a focus on identifying policies that should be given the highest priority for future research.

¹ Banerjee: MIT Department of Economics and J-PAL; Glewwe: University of Minnesota Department of Applied Economics and J-PAL; Powers: J-PAL; Wasserman: MIT Department of Economics.

² J-PAL gratefully acknowledges the funding support of the MacArthur Foundation and the Douglas B. Marshall, Jr. Family Foundation. The authors thank Amrita Ahuja, Raoul Davion, and Milena Novy-Marx for helpful comments; and Conner Brannen, Noor Iqbal, Monica Landy, and Claire Walsh for excellent research assistance.

Different countries define primary and secondary schooling differently, and in many countries students attend middle schools, upper primary schools, or junior secondary schools before attending secondary school. For the purpose of this review, we take “post-primary education” to include everything from upper primary, middle, or junior secondary school through tertiary education, as defined by the local context in different countries, including vocational school and other alternative tracks for this age group. In practice, this means that in the research we review, the majority of children are in 5th grade (i.e. 10-11 years old) and older.

The review is organized as follows. Section II provides some background on post-primary education in the developing world. Section III explains how papers were selected for this review. Section IV presents a conceptual framework for thinking about post-primary education (PPE), including a brief discussion of measuring outcomes. Section V reviews the evidence pertaining to the demand for schooling (the impact of policies that attempt to increase the willingness of households to send their children to school), and Section VI reviews the evidence on the supply of schooling (the impact of policies that change school and teacher characteristics, and more generally how schools are organized). A final section summarizes the findings, highlighting several research gaps that should receive high priority in future research.

II. Background on Post-Primary Education

There is a large body of evidence collected by economists and others showing that education increases workers’ productivity and thus increases their incomes.³ There is also evidence of the many non-monetary benefits of education, such as improved health status and lowered crime (Lochner, 2011) and, among girls, delayed age of marriage and reduced fertility (Ozier, 2010; Schultz, 2002). Inspired in part by this evidence, over the past two to three decades both governments and parents have invested heavily in their children’s education. The increases in enrollment over this period, particularly at the primary level, have been quite dramatic. From 1980 to 2010, primary and secondary enrollment rates have increased in all regions of the developing world (Table 1), so that by 2010 gross primary enrollment rates were at or above 100 percent everywhere in the world, and gross secondary enrollment rates were above 50 percent everywhere except in Sub-Saharan Africa.⁴ Moreover, Table 2 shows that primary school completion rates increased almost everywhere between 1991 and 2010, and were close to 100 percent in East Asia and Latin America, close to 90 percent in South Asia and the Middle East and North Africa, and close to 79 percent in Sub-Saharan Africa.

³ The majority of this work, following the seminal studies of Gary Becker (1964) and Jacob Mincer (1970, 1974), has focused on how school attainment relates to individual earnings, and there are now estimates of the return to schooling for many, if not most, countries in the world. More recent work has added measures of achievement to this (e.g., Mulligan (1999), Murnane et al. (2000) and Lazear (2003)), although little of this relates to developing countries (exceptions are Glewwe (1991) and Hanushek and Zhang (2009)). Recent evidence on the impact of years of education on income in developing countries is summarized in subsection IV.D.

⁴ Gross enrollment rates compare numbers of school children to the size of a specific age cohort so that grade repetition, late enrollment, and other factors can lead to gross enrollment rates over 100 percent.

Gender gaps in primary school enrollment vary widely by region, with girls at a large disadvantage in Sub-Saharan Africa, the Middle East and North Africa, and South Asia; at a slight disadvantage in Latin America; and roughly at parity with boys in East Asia and the Pacific (Table 3). In the regions with the largest gender gaps, these gaps persist (and in most places become larger, in percentage terms) at the secondary and tertiary levels. In East Asia and Latin America, by contrast, there is a gender gap in favor of girls at the secondary and tertiary levels.

This progress in access to schooling has quite naturally redirected attention towards student learning, and here the evidence on outcomes is decidedly more mixed. Over the past decade, it has become possible to follow changes in student performance in some countries on tests offered by the Programme for International Student Assessment (PISA), which assesses 15 year olds' competencies in reading, mathematics, and science. While student learning appears to be increasing in several countries, this tendency is not universal. More specifically, Table 4 presents evidence on learning among 15-year old-students in 12 countries (of which seven are in Latin America). Examining trends from 2000 to 2009, five countries show clear upward trends (Chile, Colombia, Peru, Tunisia and Turkey), while the rest show either mixed or even downward trends. In part this may simply reflect the fact that expanded enrollment brings in progressively less qualified or less advantaged students, who pull down the average score. But some of the countries with mixed or declining trends did not have large increases in school enrollment, and were increasing real educational expenditures per student. For example, in Argentina the gross secondary school enrollment rate has been about 85 percent from 1998 to 2007, and spending per pupil was somewhat higher in 2004-06 than in 1998-2000; yet test scores in 2007 were lower than in 2000. Similarly, Brazil's progress has been uneven at best, yet it experienced only a moderate increase in secondary school enrollment (7-13 percentage points) from 2000 to 2007, and real spending on education steadily increased over time.⁵

Learning lags at the secondary level are surely in part inherited from the primary schooling system. Large surveys in countries like India (ASER), Kenya, Tanzania and Uganda (UWEZO) and Ghana (TCAI baseline) suggest that a majority of the children go through the primary schooling system without mastering the basic primary level skills (ASER 2012; UWEZO 2011a-c; Innovations for Poverty Action 2011). While many countries have an entrance exam that prevents students who have failed to master primary-level skills from advancing to secondary school, many countries do not, simply taking the top students from the primary school exit exam, however unprepared they may be. It is plausible that there are also problems at the post-primary level – for example, in the quality of teaching, the pedagogical methods used, the curriculum design, and so on.

An additional challenge is the need for very large numbers of teachers capable of teaching at relatively advanced levels. This is likely to be particularly an issue in

⁵ See the World Bank's World Development Indicators. Note that Brazil's gross (net) secondary school enrollment rate increased from 99 (66) in 1999 to 106 (79) in 2005. Educational expenditures (in terms of real USD per secondary student) increased from, on average, about 1340 (350) from 1998 to 2000 to about 1510 (500) from 2004 to 2006 in Argentina (Brazil).

countries where previously only a small fraction of the population were enrolled beyond the primary level and therefore there is a very limited supply of potential teachers.⁶ To make matters worse, returns to schooling have been rising the world over and therefore these relatively more educated potential teachers are becoming more and more expensive as demand for their skills increases.

Unfortunately, there is relatively little research that can help us identify the specific deficiencies that are holding back the delivery of high-quality, relevant post-primary education and *a fortiori*, about how to make the delivery more effective. While there is an enormous body of research trying to identify the effectiveness of specific interventions into the education process, most of this research is about primary education. For example, a recent review by Glewwe et al. (2013) focused on 79 studies published between 1990 and 2010 that were considered to be of high quality. Of these studies, 51 were on primary schools while 14 were on lower secondary (middle school or junior secondary school) and only 14 were on secondary or post-secondary education. Similarly, the recent review by Kremer and Holla (2009) on randomized control trials on education in developing countries summarized results from 55 studies on primary school, 10 on lower secondary or middle school, and 13 on secondary or post-secondary education.

III. Selection of Papers for this Review

The list of papers for this study was compiled as follows. To start, the authors gathered all of the papers in Glewwe et al. (2013) that deal with post-primary education, a total of 28 papers, of which 18 were deemed by Glewwe et al. to be of high empirical quality. Those 28 papers are part of a literature review that examined research on both primary and secondary education that has done since 1990. The starting point was over 9,000 papers published since 1990 in both the education literature and the economics literature. Most did not focus on developing countries, and of those that did many did not attempt to estimate the impact of student and/or school characteristics on students' educational outcomes. See Glewwe et al. for a detailed explanation of the process by which papers were selected for that review.

An additional 10 randomized evaluations on post-primary schooling that were not in the 28 papers from Glewwe et al.'s review were obtained from the J-PAL evaluation database.⁷ The authors then obtained proceedings for past American Economic Association (AEA) and the Northeast Universities Development Consortium (NEUDC) conferences from 2005 to 2011. (The program from the 2006 NEUDC conference could not be located online and was no longer available from the host institution.) An additional 26 papers that were not already on the list were obtained from the NEUDC programs, and

⁶ A very nice illustration of this issue, in the context of primary schooling is in Andrabi, Das and Khwaja (2011) where they show that villages in Pakistan that had secondary schools for girls a generation ago are more likely to have primary schools today (run by the girls who went to those secondary schools).

⁷ The database, available at www.povertyactionlab.org/evaluations, tracks randomized evaluations related to poverty alleviation conducted by J-PAL affiliated professors. As of February 2013 there were 360 evaluations in the database conducted by 72 J-PAL affiliates.

then another 4 were obtained from the AEA programs. In addition, the authors reviewed NBER working papers dated 2005-2012, and conducted searches of World Bank publications and other economics paper databases.⁸ The citation lists of the papers assembled through the previously described strategies were also scrutinized for relevant, high-quality studies.

The authors then reviewed each of the papers from this initial pool for the quality of their identification strategy and contribution to the existing literature. In other words, the authors sought to determine which papers persuasively establish cause-and-effect relationships, as opposed to correlations;⁹ and which papers make original contributions to our knowledge of post-primary education. On this basis, a total of 56 papers were deemed appropriate for inclusion. These include studies that utilize randomized evaluation techniques and also studies that exploit quasi-random variation or natural experiments. When relevant, the authors also reference randomized evaluations on primary education, which are usefully summarized in Kremer and Holla (2009).

IV. Framework for Thinking about PPE

This section provides a conceptual framework for interpreting the research, and spotting the research gaps, on post-primary education in developing countries. The first subsection provides a general framework, based on economic theory, for understanding the education decisions of children, their parents, governments and private providers of education. The second subsection focuses on the characteristics of post-primary education that distinguish it from primary education. The third subsection discusses measurement of outcomes at the post-primary level, which in practice may help structure what types of questions are asked, and what types are omitted.

A. Conceptual Framework

Economic theory provides a useful starting point for developing a conceptual framework for understanding the decisions of households and governments regarding education. In very general terms, the amount and type of education obtained by the population is determined by both the demand for, and the supply of, education. The demand side of the process consists of households' decisions regarding the costs and

⁸ Bureau for Research and Economic Analysis of Development (BREAD) working papers, International Growth Centre (IGC) publications, World Bank policy/working papers, and issues of the American Economic Review, the Quarterly Journal of Economics, American Economic Journal: Applied Economics, Review of Economics and Statistics, and the Journal of Development Economics from 2005-2012.

⁹ A central challenge in empirical work is that the causal relationship between two variables often runs in both directions (e.g. better education may contribute to better health and vice versa), or is mediated by third factors (e.g. higher socioeconomic status may contribute both to better education and better health). J-PAL advocates the use of randomized evaluations for precisely this reason: through random assignment, the researcher can ensure the only possible explanation for differences in outcomes between treatment and comparison groups is the intervention being studied. However, other methodologies can also be useful for isolating causal relationships, and these will be briefly explained as appropriate in footnotes throughout this review.

benefits of enrolling their children in another year of school, taking the supply of education (the availability of schools, tuition and other fees, the quality of schooling, and the types of skills training offered) as given. The supply side consists of government decisions, and in many countries, private provider decisions, regarding school construction, teacher training, curriculum offered, tuition and fees, and the overall system of managing schools. The interaction of these supply and demand factors leads to education decisions that determine children's education outcomes.

On the demand side, each year parents decide whether to enroll their children in another year of schooling, which at the post-primary level could include the choice of the type of school (e.g. general secondary vs. vocational secondary). There are also ongoing choices about how much effort to put into schoolwork, how much additional teaching support (e.g. tutoring) to buy, and how much effort to put into making sure that the child does what he or she needs to do. As children become older, they play a greater role in making these decisions, and parents and children may not always see eye to eye. The basic decision rule is always the same: if the (perceived) benefits exceed the (perceived) costs for the unit taking the decision (which may be the parent, the child, or some combination of both, acting harmoniously or otherwise), and making the investment is feasible, then the investment is made. For some households, this comparison of costs and benefits may be a straightforward investment decision, where the sole benefits are future increases in income. But more generally there could be other benefits, such as a desire for education for its own sake, or for social prestige.

The demand for education at different levels (primary, secondary and post-secondary) depends on the costs, which include direct fees, the opportunity cost of the child's time (including travel time to the nearest school), and "optional" costs such as transportation fees and purchases of educational materials. It also depends on the perceived benefits, which include higher incomes, better health outcomes, social prestige, and a direct desire to be educated. Note that decisions are made based on perceived benefits, which in situations where accurate information is scarce could be very different from the actual benefits. Furthermore, the benefits from education may also depend on educational quality and the mix of skills transmitted. It is quite likely that educational benefits are not uniform across school types (e.g. private versus public, general versus vocational), locations (rural versus urban), and degree programs.

The ability and the willingness of a family to pay these costs and therefore to articulate its demand in the marketplace depends on a variety of factors such as access to credit, parental expectations about how much they will benefit from their children's additional earnings, the nature of the "contract" within the family more generally, and the community norms about education.¹⁰ There are also a range of psychological factors that need to be taken into account: For example, children often rebel against the tedium of school, even when going to school is clearly in their interest, and parents and/or teachers may or may not have the time or patience to make sure that this myopic reaction does not have long-term consequences for the child.¹¹

¹⁰ See Banerjee (2004)

¹¹ Kremer and Holla (2009)

Policies can also influence the demand for schooling. Many countries have “scholarships” for students who either perform well or simply stick it out at school. In Bangladesh, for example, the Female Stipend Program was created to promote enrollment and retention of girls in secondary school and, indirectly, to increase girls’ age at marriage and thus reduce fertility.

On the supply side, governments make a wide variety of decisions regarding how to provide education services. In most developing countries, governments are responsible for building and maintaining public schools, although in some countries local communities are expected to contribute substantial amounts for school construction while the government pays for the cost of teachers. The quality of the school infrastructure varies enormously from country to country. In almost all developing countries, governments are responsible for training and certifying teachers, and in most cases decide which teachers are allocated to which schools.

Governments must also develop a national curriculum, which involves many detailed decisions, such as the subjects to be taught, the language(s) in which they are to be taught, the level of difficulty for each grade, and methods for assessing student progress. They also establish the system by which students qualify to move on to the next grade and, more generally, to the next level of education. At the post-primary level, curriculum decisions become more complicated because there are often two or more types of secondary education (e.g. general vs. vocational) and even within these types there may be different curriculum options for students. For example, within general secondary education there could be separate tracks for math and science, arts and literature, and business and commerce, and within vocational secondary education there may be many different tracks (such as mechanical, electrical, agriculture, and healthcare).

Another aspect of supply is the tuition and other fees charged, which are generally low (often zero) at the primary level but increase as students move on to higher grades and levels of education. In many cases, households deemed to be poor are given partial or full exemptions from paying school fees, but this varies widely by country.

The government also sets rules about who can go to which school and to which classroom within the school – through its policies on admissions and tracking. This affects the peer group that each student becomes a part of, as well as the behavior of teachers and students (see Kremer and Holla, 2009 for a discussion). A classic example of this discussion is the debate over whether schools should be single-sex. Finally, school systems, both inside and outside the government, also set up processes for involving parents, both informally and more formally through parent-teacher committees.

In many developing countries private schools, both religious and secular, enroll substantial numbers of students. In 2008, the median rate of enrollment in private primary schools as a percentage of total enrollment was 10 percent in East Asia and the Pacific; 19 percent in Latin America and the Caribbean; and 9 percent in Sub-Saharan Africa. In that same year, the median rate of enrollment in private secondary schools as a

percentage of total enrollment was 16 percent in East Asia and the Pacific; 21 percent in Latin America and the Caribbean; and 15 percent in Sub-Saharan Africa (UNESCO 2011). The types of private schools and their relationships with government Ministries of Education vary widely across developing countries. In some countries, governments provide partial assistance to private schools, while in others governments provide vouchers that students can use to enroll in private schools. Regulation of private schools in terms of curriculum and certification also varies widely across countries. One general characteristic of private schools is that the costs to families to enroll their children are higher than the costs of enrolling in public schools.

There is also a wide variety of vocational education opportunities offered by the private sector, ranging from informal on-the-job training to apprenticeships to privately operated vocational schools. These also vary widely by country and data on these that are comparable across countries are quite scarce.

B. Distinguishing Aspects of Post-Primary Education

A key issue in evaluating the literature on post-primary education is the extent to which post-primary education differs from primary education. If the differences are relatively minor, the large amount of research on primary schools in developing countries would be relevant for post-primary education. For example, it is possible that some interventions that are effective at the primary level, especially in the upper grades of primary education, may also be effective in the lower grades of secondary education. Yet it is less likely that findings for primary education are applicable for upper grades of secondary education, or for vocational secondary education, and it is very unlikely that they will be applicable for post-secondary education. This section sets the stage for thinking about the applicability of results in primary education for post-primary education by highlighting the aspects of post-primary education that are different from primary education.

There are several ways in which secondary and post-secondary schools are different from primary schools. First, given that the majority of children in developing countries will enroll in secondary school upon completion of primary school, the mandate of primary schools is increasingly one of teaching fundamental skills, such as literacy and numeracy, in a way that prepares their students for the more demanding curriculum offered by secondary schools. Thus the mandate of primary schools is less and less to prepare students for immediate employment. (In some countries primary education was never designed to prepare students for employment, even though employment was the next destination for most primary school students). In contrast, the primary mandate of secondary and post-secondary education is to provide students the skills that they will need to become productive workers.

A second difference, which follows from the first, is that secondary schools are likely to be more expensive in terms of the costs of operation (teachers need to have specialized skills, laboratories and other equipment need to be provided) and unless this is entirely covered by the government it will lead to higher tuition and other school fees

charged to students. The school day is also typically longer, requiring more administrator and teacher hours. If secondary schools are funded by the government, then they show up as bigger line items in government budgets, which can be a problem in obtaining adequate funding.

Third, secondary and post-secondary schools tend to be larger in size, given the variety of courses they need to staff, which means that there are fewer of them (relative to primary schools) and so the distances students must travel to attend them are longer. In addition, governments have typically rationed secondary education. Until recently, in many countries governments have not intended to extend the option of secondary education to the entire population; rather it has been a service offered to an “elite” number of students and has focused on college preparation. Rural areas have been particularly disadvantaged in this regard.

Fourth, as discussed above, finding competent teachers at higher levels is likely to be much more of a challenge, especially in countries where relatively few people went through secondary and post-secondary schools in the previous generation.

Fifth, many successful innovations in improving skills at the primary level have been accomplished through the use of supplementary teachers who might have only a high-school education and have been given a very brief training.¹² Whether this kind of model can be made to work for teaching more advanced materials is an open question.

Sixth, there is more scope for variation in the curriculum; while primary schools focus on basic reading and math skills, secondary schools can emphasize a wide variety of curricula, such as vocational skills vs. academic preparation for post-secondary education.

Finally, some secondary schools may be elite schools that students can enroll in only if they pass an entrance exam, and in some countries entry into any school at the secondary level, especially the upper secondary level, is conditional on passing an entrance exam.

Students of secondary school age (and post-secondary school age) also differ in important ways from primary school age students. First, the opportunity costs of students’ time are higher because they are physically larger and more capable workers. Second, while primary school enrollment decisions are almost always made by parents, when a child is of secondary and post-secondary school age he or she may play a much larger role in decision making. Third, the onset of puberty can have important impacts, especially for girls. For example, the onset of menstruation could have an effect on girls’ enrollment for both biological and cultural reasons, and parents may worry about girls’ safety both when traveling to school and when in the school. At the time of adolescence the option of early marriage and demands for care of siblings and other household chores impact the opportunity cost of sending female children to school. Differences may also exist in perceived returns to educating girls versus boys. Fourth, while most parents have

¹² See Banerjee et al. 2007; Banerjee et al. 2010; Banerjee et al. 2012; He, Linden, and MacLeod 2008.

been to primary school, fewer have attended secondary and post-secondary school and so they are less familiar with post-primary education; their information may be fragmentary and in some cases very inaccurate, which is likely to increase the chance that they make poor decisions with respect to their children's education.

C. Measuring Outcomes at the Post-Primary Level

The existing literature on post-primary education primarily measures outcomes related to access and quality. Indicators of access include school enrollment, day-to-day attendance, continuation and retention, and whether students take standardized exams.

Test scores often serve as a proxy for educational quality since, at least in principle, they measure how much students have learned over a period of time. Because tests differ widely in subject-matter coverage, methodology, grading scale, and many other dimensions, most of the results in the literature are reported in terms of standard deviations.¹³ Standard deviations have the virtue of providing a common unit of comparison across different studies, but such comparisons should be made with caution.¹⁴ In the education literature, a program or policy impact of less than 0.1 standard deviations is typically considered to be a small effect, while more than 0.3 standard deviations is considered a large effect, and 0.5 standard deviations would be a very large effect. To provide a sense of absolute scale, two standard deviations is approximately the difference between the lowest-performing student in a class and the average student in the class, or between the average student and the highest-performing student. In this review, estimated impacts of education policies and programs on student learning will be compared by using the standard deviation of student test scores as the unit of measurement, keeping in mind the potential problems with such comparisons (see footnote 14).

Few of the papers in this review attempt to measure outcomes that capture the longer-term relevance of post-primary education. The best candidates for indicators of educational relevance would be labor market outcomes. These include indicator variables for employment versus unemployment, formal and informal sector employment, and self-employment; hours worked; occupation and sector; wages, benefits, and self-employment profits. Ideally studies of educational relevance would track students over many years of their working lives to capture outcomes such as long-term labor market attachment, firm

¹³ Mathematically, a standard deviation is the square root of the variance of a given set of test scores. An intuitive way to think about a standard deviation is the expected distance between a randomly selected student's score and the average score.

¹⁴ For example, two interventions may both produce a 0.3 standard deviation improvement in test scores, but this tells us little about absolute values. In a fairly homogeneous student population, for example, 0.3 standard deviation means much less in absolute terms than it would in a very heterogeneous student population (i.e. highly dispersed test scores). Policymakers may also care much more about a 0.3 standard deviation improvement in a very disadvantaged student population than in an already high-performing set of students.

survival and growth, and evolution in wages, but such tracking is difficult and expensive, so such studies are rare.¹⁵

V. Existing research: Demand for Education

Economic Theory and the Demand for Education

According to the very simple framework given above, investment in education is driven by perceived benefits and perceived costs, but may in addition be constrained by the family's ability to secure the necessary resources. The main benefits are higher incomes over the child's working life, which accrue over several decades into the future. Another benefit is better health during adulthood, both for the child and the child's own children. There may also be immediate satisfaction from education and benefits from status, as mentioned above. In the studies conducted to date, researchers have generally focused on measuring the income benefits, and in some cases the health benefits, associated with education. In these studies, the costs of education have been tuition and other fees, both required and optional, and the opportunity cost of the child's time.

According to the basic economic model, if there are no credit constraints, and the income benefits are perceived to be much larger than the other benefits, the schooling decision is a relatively simple investment decision, where investment continues until the point is reached where the increase in the present discounted value of the future income stream (which depends on the rate of interest) from more investment is outweighed by the cost. This simple model of the schooling decision has several immediate implications. First, if the costs of schooling (which include borrowing and transportation costs) increase, the optimal investment in schooling decreases. Second, government policies that directly raise the return to schooling (such as scholarships) will encourage students to put more effort into their education (including staying longer in school), though the effect will vary by the students' perceived probability of winning the award. Third, if the wages paid to educated workers increases, the optimal level of schooling increases and conversely, if the wages paid to workers with little or no education increases (that is, the opportunity cost of sending a child to school increases), the optimal investment in schooling goes down. Fourth if there is reason to expect that a child will not be able to make use of her education (say because of gender bias or some other form of discrimination) there will be less investment in education. Fifth, if the quality of schooling increases, so that students obtain more productive skills for a given number of years in school, their expected wages will also increase and thus the optimal investment in schooling will increase.

Another interesting (and perhaps counterintuitive) implication of this simple model is that household income, and all other household attributes other than the aptitude of the child (and perhaps the parents' ability to provide homework help) should not affect how long, how hard, or what subject the child studies. However, there are several reasons

¹⁵ For an illustration of a paper that follows a sample over a decade and includes a large set of labor market outcome indicators, see Baird et al. (2012).

why this may not be the case. First, if schooling is seen as intrinsically valuable or if it conveys benefits in terms of status – so that it is no longer merely an investment – wealthier families will choose higher levels of schooling than otherwise identical but poorer families. Second, the implicit assumption behind this result is that *the child* can borrow as much as she needs to invest in her education. If, as is more plausible, parents are the ones who can borrow, then the level of investment depends on to what extent the child's increase in earnings compensates the parents (he or she could pay them back either in money or in old age care, or they could directly feel compensated by the increase in his or her well-being). Third, if most parents themselves are credit constrained, less wealthy households may not be able to pay the costs of schooling even if the benefits outweigh the costs, and so they will be unable to provide their children the optimal years of schooling while wealthy households are more likely to be able to do so. There would be a similar distortion if the family is savings constrained, in the sense of not being able to hold on to its earnings until it is time to pay school fees. Fourth, it is possible that some families are often unable to act on incentives despite the clear benefits of doing so: self-control problems and intra-family conflict are some reasons why this might be the case. The child, being a child, may prefer to play rather than do homework; parents, being busy or otherwise distracted, might not put enough effort into making sure that the homework gets done; or there may be a collective action problem between the parents, where each parent “free-rides” on the other parent monitoring the child. These problems may be more severe for poorer families, because they might be under greater stress for other reasons. Finally, it may be the case that poorer households have less accurate information that leads them to underestimate the wage gains from their child attaining a higher level of education. In this case less wealthy households will choose lower levels of schooling than wealthier households due to this inaccurate information.

If there is underinvestment in the education of poorer children because parents do not necessarily value the future earnings of children, or because of credit constraints or savings constraints, there is a clear argument for government intervention. The obvious response is to address the source of the problem directly by reducing the price or providing better opportunities to borrow or save. Alternately, one could try to put more resources into the hands of families by providing them with income support. If the problem arises from commitment issues, the policy response may be to provide some relatively immediate incentives based on either attendance or performance in school (immediate, because the problem is that people who ignore long term benefits end up acting myopically). Whether it is optimal to give the incentive to the child or to his or her parents depends, of course, on whose commitment one cares about, as well as the nature of the relations within the family. On the other hand, if the problem is lack of information, the natural response is to provide information. Information provision can take the form of notifying parents and children of the true benefits of education, and of the differing quality of accessible educational institutions. Finally, if the difference between rich and poor children reflects the consumption benefits of education, the policy response is less obvious. On one hand, when all parents fully internalize the value of the returns on educational investment that their children obtain, and there are no commitment issues, the consumption value of education leads to over-investment, and poorer children are actually more likely to get the right level of investment. However, without

discounting the possibility that some wealthy parents do over-invest, underinvestment seems *prima facie* the more likely option, at least as far as the poor are concerned.

With this simple theoretical framework we now turn to the evidence on policies that affect the demand for schooling. We start with price-based policies, including conditional cash transfers and other incentives, followed by income support policies and informational policies.¹⁶

A. Price-Based Policies

Price-based policies can take two forms: unconditional and conditional price changes. An unconditional price change is a change in the cost of schooling that does not depend on any family or individual behavior. For example, a reduction in school tuition from \$10 per year to \$8 per year would be an unconditional price reduction. As suggested above, a reduction in the price of obtaining an education both raises the rate of return for each additional year of education and makes education more affordable for households facing credit or other constraints. Price changes that go into effect only if a household or an individual engages in a certain behavior (such as regular school attendance) are known as conditional price changes. Such price-based policies (e.g. the price is lower only for those who attend school regularly, for example) may help deal with commitment issues. The most common of the conditional price changes is a conditional cash transfer (CCT), which provides regular cash payments to students or the students' parents if the student satisfies an attendance requirement. Price reductions, either conditional or unconditional, should increase time in school. A price increase, on the other hand, should have the opposite effect. What theory cannot tell us, however, is the size of any of these effects, i.e. how responsive school attendance or educational attainment is to price changes. If the effect of a price increase is small, it may be optimal to raise tuition and use the funds raised to improve school quality, the net effect of which could be to increase enrollment. On the other hand, if the effect is large an increase in tuition could have a large negative impact on enrollment and may produce little revenue with which to improve school quality. Policymakers may also be concerned about which students are selected into and out of education by price changes, i.e. if poorer students are selected out by a price increase, this has very different implications than if those with the lowest returns to education were selected out.

Unconditional price reductions

The most obvious price based policies are unconditional price cuts (or unconditional price increases). There are a few recent studies that look at the impact on enrollment of unconditional price reductions. The first, by Barrera-Osorio, Linden, and

¹⁶ We decided to not cover policies that influence the human capital development of primary age or even younger children that might subsequently affect their demand for post-primary education in this survey. The obvious examples of such policies include primary education policies, as well as school health and early childhood development policies. While recognizing that these policies can be enormously important, we limit the discussion here to post-primary education policies, which are the main focus of this paper.

Urquiola (2007), finds mixed results from the *Gratuidad* initiative, a proxy means-tested school fee reduction program in Bogota, Colombia. Monthly educational spending for high school students in this population ranged from \$11 to \$29, or between about 10 and 25 percent of the monthly minimum wage. The program offered school fee reductions to children in the bottom two of six categories based on a proxy-mean index called *Sisben*. The authors use a regression discontinuity design that exploits the discrete changes in school fee reductions around the cutoff scores for the lowest two *Sisben* categories.¹⁷ The authors find a 6 percentage point increase in high school enrollment among students just past the threshold for the second-poorest category (*Sisben* 2), who experienced an approximately 50 percent fee reduction, relative to those on the other side of the cutoff with no fee reduction. However, there was no statistically significant difference between students on either side of the cutoff between *Sisben* 2 and *Sisben* 1, which had a 100 percent fee reduction. The authors also find some unexplained heterogeneity by gender: at the secondary level, the enrollment increase was driven primarily by girls, while at the primary level it was driven by boys.

Borkum (2012) combines fixed effects and regression discontinuity analysis to examine a targeted fee-elimination initiative in South Africa. Primary and secondary schools were divided into national quintiles based on the poverty score of the surrounding areas, and the poorest two quintiles were required to eliminate school fees starting in 2007. The median secondary school fee was R130 (\$49 PPP) per annum, a modest 1.5% of the median household's annual income in fee elimination eligible areas. The fixed effects regressions for South Africa as a whole suggest that the program increased national secondary school enrollment by almost 2 percentage points, with effects concentrated in the poorest quintile of schools and in the earlier secondary school grades (8-10).

The paper also tries to use a regression discontinuity approach to get at the same question, using the schools at the threshold between the second and third quintiles in the Eastern Cape Province, which implemented the policy strictly according to the poverty index cutoffs, unlike other provinces that made discretionary adjustments. Estimates from the RD analysis are close to zero and imprecisely estimated. However, fixed effects analysis on the same data from Eastern Cape Province finds an average effect of over 3 percentage points, and an effect of almost 6 percentage points within the poorest quintile. Taken together, these results suggest that the fee reduction program had the largest impact on enrollment in the poorest areas, especially given that the fees did not represent a large amount of money for the households concerned, with perhaps relatively little effect near the cutoff for fee elimination (for those families who are wealthier and therefore can more easily pay the fees).

¹⁷ Regression discontinuity (RD) is a method for estimating causal relationships between variables in the absence of a controlled experiment. The technique exploits a sharp cutoff in some continuous variable (in this case, the *Sisben* index) that determines a discontinuous policy response (in this case a 0 versus a 50 percent fee reduction, or a 50 versus a 100 percent fee reduction). Assuming that the eligibility variable is hard for people to manipulate precisely, people on either side of the cutoff are likely to be very similar, so differences in outcomes (enrollment, in this case) can be attributed to the intervention. One weakness of this approach is that the causal estimates may be unbiased only for people who are close to the cutoff.

Spohr (2003) analyzes the effects of a 1968 policy reform in Taiwan that extended tuition-free, compulsory schooling from 6 to 9 years, through the end of junior secondary school. He estimates that the reform led to an average increase of 0.4 years of schooling for males, and 0.25 years for females. Using birth after 1954 as an instrument for years of completed schooling, 2SLS regressions suggest that each year of schooling increased the likelihood of paid employment by 2 percentage points for males and over 5 percentage points for females.¹⁸

Wydick, Glewwe and Rutledge (2013) examine the Compassion International child sponsorship program using data from six countries (Bolivia, Guatemala, India, Kenya, the Philippines, and Uganda). Individual sponsors in developed countries sponsor individual children in developing countries, paying \$25 - \$40 per month that is used to pay for students' school fees, uniforms and school supplies, as well as for tutoring, health care and group activities that include religious instruction. Sponsors typically provide support for many years, until the child finishes secondary school, and most sponsors regularly send children letters as well as birthday and Christmas presents. Using a variety of different estimators (household fixed effects, regression discontinuity, and IV methods), the authors find that the program increases years in schooling by 1.0 to 1.5 years, and also increases the probability of obtaining a white collar job. In five of the six countries, they find greater impacts for the gender that has lower levels of schooling among untreated children: girls in the case of Uganda, Guatemala, and Bolivia; and boys in the case of India and the Philippines. In Kenya, the impact of sponsorship is somewhat large for boys even though boys have slightly higher initial levels of schooling.

Andrabi, Das, and Khwaja (2011) find evidence from Pakistan that a lack of qualified teachers is a constraint to the provision of education, and that investments in secondary schools can foster the growth of private schools by increasing the local supply of potential teachers. They use instrumental variable analysis to show that villages where girls' secondary schools (GSS) were constructed were 27 percentage points more likely to see private primary schools emerge in the following years, on a base of 12 percent. Identification is based on eligibility guidelines for new GSS in the 1980s, which gave preference to villages with higher populations that did not have another GSS within a 10-kilometer radius. To show the case for interpreting this causally, they show that the eligibility status was not correlated with any other observable socio-demographic characteristics and that the GSS eligibility rule does not predict any other type of public investment, including other types of public schools.

¹⁸ Instrumental variables (IV) are an econometric technique that estimates the causal relationship between two variables when there is some difficulty disentangling the direction of causality or ruling out the influence of third factors. In the case of employment and schooling, for example, an individual's cognitive ability is likely to increase both her educational attainment and her likelihood of paid employment, so measuring a simple correlation between schooling and employment is likely to be biased—some of the apparent effect of education on employment may actually be explained by more educated people also having higher cognitive ability. The IV technique relies on finding an “instrument” whose only effect on the outcome of interest (employment) is through its effect on the hypothesized cause (schooling). In the case of Spohr's (2003) study, the first stage of the analysis relates the instrument (the extension of free, compulsory schooling) to years of schooling, and the second stage relates the change in schooling caused by the policy to subsequent changes in employment rates.

The authors also argue that the effects of GSS construction on private schools operate through a supply channel and not solely through demand channels. (An example of the latter would be if better-educated mothers were more likely to send their children to private schools.) They show that the presence of a GSS more than doubles the stock of educated women in the median village and reduces the average local teaching wage by 27 percent. If the effects of GSS operated solely through demand channels, we would expect an increase, rather than a decrease, in teacher wages.

Conditional price reductions

Conditional cash transfers (CCTs) add to price reductions an incentive component: these programs provide regular (usually monthly) payments to parents if their children are enrolled in school (often with an 80-90% attendance requirement). In effect, these programs amount to a subsidy for attending school. The following paragraphs summarize several recent CCT studies, almost all of which were implemented as randomized control trials (RCTs).

Schultz (2004) evaluated the first large-scale CCT program, Mexico's *PROGRESA* program. In 1997, the Mexican government launched *PROGRESA* to improve health and education outcomes among the poor. The cash transfer benefit from *PROGRESA* (now called *Oportunidades*) comes in two forms. The first is a monthly fixed stipend of 90 pesos (approximately US\$7) conditional on family members obtaining preventive medical care. The second type of transfer comes in the form of educational scholarships, which are given to families of children from third grade onwards, conditional on those children attending school a minimum of 85 percent of the time and not repeating a grade more than twice. Beneficiary children also receive money for school supplies once or twice per year. The size of the education stipend is larger at higher grades and is also higher for girls because the government wished to encourage girls, in particular older girls, to stay in school. The stipends vary from 60 pesos per month for children enrolled in third grade to 225 pesos per month for females enrolled in the third year of lower secondary school. The high end of the transfer range amounts to about 44 percent of the typical monthly wage for a male day laborer in the region and roughly two thirds of what the child could earn by working full time if he or she were able to find employment.

The program was implemented in 314 localities randomly chosen from among an initial group of 495 localities. Comparisons of *PROGRESA* and non-*PROGRESA* localities suggest that the program had a significant positive impact on enrollment rates. On average, the program increased years of schooling by 0.66 above a baseline level of 6.80 years (a 9.7 percent increase). The largest difference in enrollment was for those children who had completed grade 6, and were thus qualified to enroll in lower secondary school. Among this cohort, the enrollment rate increased by 11.1 percentage points. This impact was also disproportionately concentrated among girls, whose enrollment rate increased by 14.8 percentage points, compared to an increase of 6.5 percentage points for boys.

Attanasio et al. (2011) also find that *PROGRESA* had a positive effect on the enrollment of children, especially after primary school. The program had no impact on enrollment of 10 year olds, because almost all children below grade 6 already attended school, however, it increased enrollment for 14 year-old boys by 14 percentage points. The program impact increased with age as the probability of dropping out, absent the program, increased.

Todd and Wolpin (2006) develop a structural model of household decisions regarding fertility and demand for education, and use experimental data from *PROGRESA* to conduct simulations of counterfactual policies. They compare the school attendance and completion rates from the existing *PROGRESA* subsidy schedule to those from a series of hypothetical alternative subsidies. Based on these counterfactual policy experiments, they predict that eliminating the subsidy in lower grades, where school attendance is already very high, and increasing it at higher grade levels will significantly increase the proportion of children who complete nine or more years of education without increasing program costs per family.

Similar impacts of CCT programs on school enrollment have been found in Brazil, Cambodia, Colombia, Ecuador, and Pakistan. Glewwe and Kassouf (2012) find that Brazil's *Bolsa Escola/Familia* program increased enrollment by about 5.5 percent in grades 5-8. Dropout rates decreased by 0.4 percentage points in those grades, and grade promotion rates increased by 0.3 percentage points. The program was more effective at increasing female enrollment than male enrollment in grades 5-8, but there were no gender differences in the effects on dropout or promotion rates. A CCT program in Cambodia increased school enrollment and attendance in secondary schools by approximately 25 percentage points (Filmer and Schady 2009). This relatively large effect may in part be due to targeting the scholarships toward the students most likely to drop out, who were disproportionately female students from very poor families. In Colombia, Attanasio et al. (2010), utilizing a difference-in-difference analysis of the *Familias en Accion* program, found that school enrollment rates of 14-17 year olds increased by 5-7 percentage points. Schady and Araujo (2006) find that the *Bono de Desarrollo Humano* (BDH) program increased school enrollment in Ecuador by 10 percentage points. Finally, Chaudhury and Parajuli (2006) find that the Female Secondary School Stipend (FSSS) program in Punjab, Pakistan, increased enrollment among 10 to 14 year-old girls by 9 percent, and increased daily attendance rates among this same group by 10 to 13 percentage points.

Recent evidence suggests that the positive impacts of CCT programs on enrollment and attendance are sustained over the long term. Behrman et al. (2011) find that five and a half years after the start of *PROGRESA*, boys aged 9-12 before the program started achieved 0.9 to 1.0 grades of additional schooling, and the same age-cohort of girls achieved 0.7 to 0.8 additional grades, compared with similar children not receiving the program. The program also increased the schooling of boys aged 13-15 by about half a year, while it had no significant impact for girls of the same age group. Baez and Camacho (2011) find that children who participated up to nine years in a CCT

program in Colombia were 4 to 8 percentage points more likely than nonparticipant children to finish high school. These impacts were larger for girls than they were for boys.

While CCT programs have, in general, led to increased enrollment and higher school attainment, their impact on student learning levels is more uncertain. Despite the increase in enrollment, Filmer and Schady (2009a) in Cambodia, found no impact on mathematics and vocabulary tests. They suggest that this is due to the entry of lower-ability students into the school in response to the program. In the long-term, both Behrman et al. (2011) in Mexico and Baez and Camacho (2011) in Colombia also found no program impact on students' performance on achievement tests. However, of those who did graduate from high school in Baez and Camacho's (2011) study, program participants and nonparticipants performed similarly—suggesting that there is no difference in terms of “innate talent” between those children who complete school because of the program and those who would have done so in any case.

As mentioned above, CCTs effectively combine an income transfer with an incentive for school attendance and/or performance. It is useful and interesting to try to identify the relative magnitudes of these two effects. Baird, McIntosh, and Özler (2011) conducted a randomized evaluation in rural Malawi to assess the importance of the conditionality in a cash transfer program, as well as the sensitivity of the impact to the amount of the transfer. Their target population is girls between the ages of 13 and 22 who reported being enrolled in school at the time of the baseline survey. One hundred seventy-six communities were randomly assigned to treatment or control status. A subgroup of the treatment communities was then randomly assigned to receive offers for monthly cash transfers *conditional* on attending school regularly (CCT) while another group received offers for *unconditional* cash transfers (UCTs), so that they did not have to attend school regularly to receive the transfer. In the CCT intervention, the amounts of the transfers were randomly varied across communities. The transfers to parents were in two-dollar increments between \$4 and \$10. In addition, within each treatment community (both UCT and CCT), there were randomly allocated transfers to girls (rather than their parents), which were in one-dollar increments between \$1 and \$5 per month.

The average enrollment gains in the CCT intervention were significantly larger than those in the UCT intervention. Over the course of two school years (six terms), girls' enrollment in secondary schools in the UCT intervention increased by an average of 0.23 terms. However, in the CCT intervention, enrollment increased by 0.54 terms.¹⁹ The CCT program also significantly increased school attendance, and modestly increased test scores for reading, mathematics, and cognitive ability by 0.12 to 0.14 standard deviations. Moreover, after the cash payments ended, the impact of the CCT intervention appeared to persist into the following term, which suggests that increased investment in schooling

¹⁹ These findings utilize teacher-reported enrollment rates. The findings are the reverse when analyzing self-reported enrollment, though the authors present evidence that self-reported enrollment, particularly in the UCT and control groups, is subject to positive bias, i.e. a higher percentage of girls in the control and UCT groups reported that they were enrolled in school when they in fact were not. See the paper for further details.

also makes future investments more productive. No statistically significant gains in attendance or test scores were found in the UCT treatment arm. While the positive findings for student achievement in the CCT arm stand in contrast to the null findings from the other CCT papers, the authors do not offer an explanation for this key difference. It therefore remains an open question whether CCTs are also an effective means to improve students' test scores.

The importance of conditionality is also a message of another of the papers previously discussed: Schady and Araujo (2006) suggest that the fact that some households believed that there was a school enrollment requirement attached to the transfers in Ecuador, even though such a requirement was never enforced or monitored, helps explain the magnitude of the program effects on enrollment.

Another interesting finding from Baird, McIntosh and Özler (2011) is that when comparing the different levels of payment in the CCTs, the researchers found that even the smallest tested transfer of around \$5 per month (combining the minimum \$4 transfer to the parents and \$1 transfer to the individual) prompted the same increase in schooling as transfers twice as large, and the improvement in schooling outcomes was not sensitive to the identity of the CCT recipient (i.e. whether it was the girl or her parents). Filmer and Schady (2009b), mentioned above, also find evidence of significant diminishing returns to transfer size, despite the fact that even the larger transfer represented on average only 3 percent of the consumption of the median recipient households.

In a randomized experiment in Bogota, Colombia, Barrera-Orsorio, Bertrand, Linden, Perez-Calle (2011) evaluated two variations on the traditional CCT design that attempt to provide incentives not only for regular attendance, but also for re-enrollment at the start of the next school year and graduation and matriculation to tertiary (post-secondary) school. The alternative designs were also intended to address financial challenges other than day-to-day cash constraints, such as difficulty saving money for annual expenses at the start of the school year.

In the "Standard CCT" design, students received the equivalent of approximately \$15 per month if they maintained 80 percent school attendance. In the "Savings CCT" design, the monthly transfer was only \$10 with the same conditionality, but students received about \$50 at enrollment time for the next school year, so that the total transfer was roughly equivalent to the Standard CCT. In the "Graduation CCT" design, the monthly transfers were again \$10 per month, and students received roughly \$300 if they graduated from secondary school—immediately upon graduation if they provided proof of enrollment in higher education, and after a one-year delay otherwise. The first two designs were implemented in one locality on a sample of students in grades 6-11, and the third, due to its high cost, was implemented on a smaller sample of students in grades 9-11 in another locality. For that reason the impact estimates for the Graduation CCT are not directly comparable to the others, but the authors provide non-experimental comparisons between them.

The two non-standard designs appeared no less effective at increasing attendance than the standard design, despite the smaller monthly transfers, and they significantly increased re-enrollment and matriculation in tertiary school. The Standard CCT and Savings CCT increased monthly attendance by 3.3 and 2.9 percentage points, respectively, and the Graduation CCT increased attendance by 5.2 percentage points relative to its own control group. The non-standard designs had a larger impact on re-enrollment in school the following year: the standard, savings, and graduation incentives caused increases of 1.7, 4.5, and 4.2 percentage points, respectively. Finally, only the non-traditional designs had an impact on enrollment in higher education. For the students in the Graduation CCT there was a 49 percentage-point increase in matriculation (from a control group average of 19 percent enrollment), and even though it provided no direct incentive to continue, the Savings CCT also led to an increase in enrollment of 9.4 percentage points.²⁰ Notably, the differential effects of the standard and non-standard CCT arms were driven primarily by the lowest income students and those with the lowest participation rates. In the Standard and Savings CCTs, the gains in attendance experienced by girls were much smaller than those experienced by boys, perhaps because girls had higher attendance in the absence of the program than boys.²¹ The effect on enrollment, however, was approximately the same across treatments for boys and girls.

An unintended consequence of the Bogota CCT program was increased inequality in educational attainment within households. Parents could enroll any number of their children in the lottery for the program, but the average family only entered 1.3 children, compared to the average of 2.5 eligible children. This suggests that either parents have preferences about which children to educate, or perceived educational returns to be convex,²² or there was some constraint on how effectively the family could pool resources from different sources. There was also evidence that parents diverted resources toward children who won the CCT and away from their siblings, particularly girl siblings. Sisters of children who received the transfers were 10.4 percentage points less likely to be enrolled in school than sisters of children not in the program.

²⁰ The 49 percentage point increase in tertiary enrollment due to the Graduation CCT is acknowledged by Barrera-Osorio et al. to be extremely high. In the paper, they provide a discussion regarding treated respondents' incentive to misreport and whether the reported enrollment rates are consistent with secondary enrollment and graduation rates. They conclude that the estimated effect is consistent with respondents' survey responses regarding secondary graduation rates and time spent studying, as well as consistent with administrative data collected on secondary enrollment and graduation rates and enrollment.

²¹ The difference in treatment effects by gender is 3.4 percentage points for the basic treatment (significant at the 10 percent level) and 4.5 percentage points for the savings treatment (significant at the one percent level). Girls have a 3.2 percentage point higher attendance rate than boys absent the treatments, so the net effect of the treatments is to increase the attendance rate of boys to roughly that of girls.

²² Lottery entrants were more likely to be boys than girls, particularly at higher grade levels. The paper does not address why most families chose to enter only one child in the lottery. This phenomenon, paired with the subsequent analysis of households' allocation of resources favoring children who received the CCT, may indicate parents' belief in convex returns to education, i.e. for each additional year of education received, the earnings gain is progressively larger. For example, if a household is considering educating two children and intends to maximize the household's total present discounted value of earnings, then in the case of convex returns to education, the household will prefer that one child has the highest educational attainment possible. In the case of concave returns to education, the household will prefer that both children have equal levels of educational attainment.

Discussion

The body of evidence on price reductions produces a general conclusion: students' school attendance and completion are sensitive to the price of schooling. Indeed, every price reduction study reviewed in this section – whether it entails a conditional or unconditional intervention – finds positive impacts of price reductions on students' school attendance and/or completion. Furthermore, many of these studies find that the poorest households are most responsive to cost alleviation, suggesting that not only price effects but also income effects are operational, and the most disadvantaged populations likely face other types of barriers such as credit, informational and savings constraints. Students who have lower baseline levels of access to school—frequently girls—also appear to be more responsive to these types of interventions. Despite the overarching positive results of price-based policies in increasing school enrollment and attendance, the evidence on the effects of price reductions on student *performance* is less conclusive. A leading explanation for the null findings in student achievement involves the changing composition of the student body due to the program.

Open Questions

1. How responsive is student attendance to the price of schooling? Does it vary across student sub-groups? Does it depend on the students' grade?
2. How should CCT programs be designed? Should the transfer be comprised of prize/scholarship at the end of the year, or multiple short-term prizes? Does the structure of the prize affect the stronger and weaker students differentially? How about using some of the money as an incentive to matriculate or go to college, or forcing students to save a part of it which could help them finance their future education?
3. Does the identity of the transfer recipient matter? Should these transfers go to the family or be given in a way that only the student can access?
4. Is it possible to design CCTs to work on a large scale in government systems (as in India) where there is no public exam until 12th grade and attendance statistics are routinely exaggerated, and so there is no reliable information on which to condition transfers? In systems where monitoring is weak, could UCTs be more effective?
5. What are the effects of CCTs on student performance as measured by test scores, and does this vary over types of students?

B. Performance-Based Incentives

There are only a few studies on the impact of financial incentives conditional on academic performance in developing countries. In theory, these transfers act as a price subsidy on effort, potentially leading to greater effort and better outcomes.

In a randomized experiment in western Kenya, Kremer, Miguel and Thornton (2009) evaluate a school-level scholarship program for 6th grade girls. The scholarship

paid school fees for the next two academic years and a provided grant for school-related expenses for girls who scored in the top 15 percent of their district on an academic exam. Sixth grade girls in treatment schools scored between 0.12 and 0.27 standard deviations higher than girls in control schools. One of the more surprising results is that test score gains are not concentrated in the upper part of the skill distribution or among students from the wealthiest families: students who seemed to have little or no chance of qualifying for the scholarship appear to have benefitted from the introduction of this scholarship program. There is also limited evidence that boys' test scores increased as a result of the intervention. While the paper is not able to pin down the precise mechanism through which test scores – particularly for groups with little to no chance of winning the scholarship – improved, there is some evidence that teacher attendance increased in treated schools. Enhanced teacher effort and peer effects are the likely proximate causes for the positive externalities of this program.

Incentives for students could be implemented at either the individual level or the group level. Blimpo (2010) investigates the differences in the impact of individual versus team incentives in a randomized experiment in Benin. The study tracked the performance of 1,476 tenth-grade students from 63 private and 37 public secondary schools on the standardized secondary school certification examination. In the “individual target” group, students were promised \$10 (equivalent to about 4.5 weeks of the students' average reported pocket money) if they passed the exam, which typically had a pass rate slightly below 50 percent. In the “team target” group, students were randomly assigned to groups of four, and received \$40 for the team if its average score equaled or exceeded the passing score. In both individual and team targets, the reward was tripled if the student's or team average score met or exceeded the grade required for honors. In the “team tournament” group, the three teams with the highest average scores each won a \$640 prize.

The individual and team target incentives had comparable effects on test scores—0.29 and 0.27 standard deviations, respectively—though the latter was only marginally significant. The team tournament incentive had the largest effect at 0.34 standard deviations. In line with theoretical predictions, the individual incentive had the largest impact at the median (0.67), close to the pass/fail threshold, and did not have a statistically significant effect among higher-performing students at the 85th percentile. In addition to having the highest average treatment effect, the team tournament was the only incentive with significant treatment effects at all levels of student performance. Blimpo argues that the tournament incentive may have induced all teams to work harder, since the lack of a fixed target may mitigate incentives for free-riding.

Discussion

The two studies detailed above suggest an important and encouraging role for financial incentives in improving academic performance, both for students who are targeted by the incentives and – more surprisingly – for their untargeted peers. Further research and experimentation is necessary to shed light on the efficacy of the various forms of performance incentives, their longer term impact, whether their impact varies

across different sub-groups of students, and whether there is any evidence that these incentives lead to crowd-out of intrinsic effort.²³

Open Questions

1. Do non-monetary performance incentives affect student performance?
2. Do performance incentives crowd out students' intrinsic effort?
3. What are the long term effects (if any) of performance incentives?
4. Do all students respond to performance incentives or are there particular groups that are particularly responsive?
5. Are individual or group based incentives more effective? Does their relative effectiveness depend on specific circumstances?
6. Are small incentives sufficient to obtain large effects? More generally, does the marginal (per dollar) impact of incentives decline rapidly as the monetary value of the incentives increases?

C. Credit-Based Interventions

Since the poorest households are often the most responsive to price-based interventions, it is likely that these households are facing some type of financial constraint in their educational decision-making. The basic model of investments in human capital implies that households that are credit constrained will underinvest in their children's education. The research in this domain is sparse and the existing studies focus primarily on credit constraints in higher education.

Kaufmann (2012) uses data from Mexico on 15-25 year old individuals' subjective expectations for future earnings under a number of scenarios for education investment in order to shed light on the decision to attend college, conditional on high school completion. The results of the analysis suggest that the poorest students require higher expected returns in order to be induced to attend college, suggesting that these students face higher direct and indirect costs associated with college enrollment. Kaufmann then tests the sensitivity of the educational decisions of students from different wealth backgrounds to changes in direct costs of education such as tuition and living expenses. While individuals in the full sample are not responsive to changes in tuition expenses, those individuals from the poorest backgrounds who report high expected returns to college, display substantial sensitivity. This result suggests a role for credit constraints, but does not rule out other explanations such as informational barriers.

Solis (2011) uses a regression discontinuity design to examine the impact of access to credit for tertiary education in Chile. He examines two programs in Chile that

²³ "Crowd-out of intrinsic effort" refers to the phenomenon in which the introduction of a program or tangible incentive intended to spur greater effort actually results in the displacement of existing intrinsic effort. For example, the introduction of a monetary incentive associated with obtaining good grades in school might crowd out students' intrinsic effort if it diminishes students' existing motivation to obtain good grades.

give tuition loans to students in the four lowest income quintiles who score above a cutoff in the national college admissions test. The sample includes 3,438 first-time test takers (to avoid possible self-selection from students learning to take the test) between 2006 and 2009. Solis estimates that access to loans induces a 21 percentage-point increase in college enrollment. To provide a benchmark, students of comparable socioeconomic background but slightly below the test score cutoff had enrollment rates of 16 percent. The impact is largest among the poorest students. In fact, the enrollment gap between the highest and lowest income quintiles is completely eliminated for those who qualify for tuition loans, i.e. their test scores are above the cutoff, while the enrollment gap for students below the test score cutoff is similar to what is found in the general population. Eligible students also experienced an increase in their probability of re-enrollment by 33 and 29 percentage points for the second and third year of college, respectively. To test whether the RD results might generalize to other parts of the admissions test distribution, Solis identifies all twin pairs in the sample and estimates family fixed effects regressions. He reports effects of similar magnitude from access to credit in this second identification strategy.

Solis also attempts to decompose the effects on enrollment into a price effect—i.e. the impact of below-market interest rates and weak enforcement of loan repayment under the program—and an “access effect,” caused by lack of access to credit markets in the absence of the program. He shows that the effects of access to loans is similar for the bottom three income quintiles, while in the fourth quintile—which had access to private, but more expensive loans on similar terms to the government program—had a very weak enrollment discontinuity around the cutoff. This suggests that the impact of the credit program is driven primarily by the access effect rather than the price effect.

Gurgand, Lorenceau, and Melonio (2011) also use a regression discontinuity design to compare South African students above and below a credit score threshold for university loans. They collect data on nearly 10,000 loan applications to Eduloan, a private company supported by international donors, between 2004 and 2007. The sample comes mostly from middle-class South African households and the applicants are relatively old, averaging over 27 years, because Eduloan rules require regular employment. Therefore, the sample represents more employees who are looking to upgrade their skills than parents borrowing for their children’s education.

They find an impact similar in absolute magnitude to that found by Solis in Chile: loan access increased the probability of enrollment by 22 to 25 percentage points, representing in this case a 50 percent increase over the baseline enrollment rate. The impact of the loans was twice as large for the lowest income quartile.

Finally, as discussed above, Barrera-Osorio, Bertrand, Linden, Perez-Calle (2011) studied the effect of a CCT program that forced the beneficiaries to save a part of the transfer until the next enrollment season. This forced saving may be seen as a substitute for credit. This version of the program in fact had the largest effect on re-enrolment rates and also had a significant positive effect on enrolment in higher education.

Discussion

The previous four studies provide some suggestive evidence that individuals are credit constrained in their decisions to pursue higher education: individuals pursue higher education at much higher rates when given access to education loans. The fact that these loans often have interest rates that are lower than the market rate means that two interventions are simultaneously occurring: expanding credit access and lowering the cost of education due to lower interest rates. Although difficult to disentangle, there is some evidence that most of the positive effect of these loan expansion programs runs through the credit access channel.

Open Questions

1. Are households credit constrained in their educational decisions at the middle and secondary school level?
2. Is the main issue with credit constraints mainly an issue of access, with the precise interest rate being less important, or is it that there are loans available but the interest rates are very high (so that subsidies would be the appropriate policy)?
3. What policies can be implemented so that lenders can be repaid once the borrower has a job? Can the government garnish wages from government workers, or even from workers in the formal private sector?

D. Income Supplementation

Given that households seem to be credit constrained, an alternative, though less well targeted, way to get at the same problem is to provide income support to the families.

Edmonds (2006) examines the effects of a policy experiment in South Africa—the extension of the public Old Age Pension (OAP) to black South Africans after the fall of apartheid—on schooling decisions. The transfers under the OAP were quite large, 125 percent of black median per capita income. This setting is well suited to test the model for several reasons. First, the pensions were highly anticipated by recipient households: they were linked to age (60 for women, 65 for men), the means test was binding for very few black households, and the program was well known. Second, black South African households are typically multi-generational, and other research has shown that the OAP income is widely shared with other family members.

Edmonds' sample includes school-age children in 3,708 rural households who co-reside with an elder between the ages of 50 and 75. The data come from the June 1999 Survey of the Activities of Youth in South Africa (SAYP). Identification of the effect of pension eligibility is achieved by regressing outcome variables (school enrollment and attainment, child labor) on indicator variables for the presence of pension-eligible male

and female elders. The estimates also control for the age of elder household members to account for possible age trends and differences in expected permanent income. Therefore, the coefficients on pension eligibility compare pension-eligible households to nearly eligible households with approximately the same expected permanent income.

Edmonds finds a significant, 18 percentage-point increase in school enrollment for 13- to 17-year-old boys who live with pension-eligible males. There were no significant effects of pensions on girls' enrollment, nor was there a significant effect from female pensioners. It is important to note that boys had lower enrollment rates to begin with, and that boys with nearly-eligible females had higher enrollment than boys with nearly-eligible males, thus the data suggest that pensions helped boys living with male elders to "catch up" rather than surpass other children's enrollment. Similarly, in the schooling attainment estimates each additional year of exposure to a male pensioner increased attainment for boys, but there were no effects for girls or for female pension eligibility. Edmonds hypothesizes that men may be more likely to be credit-constrained, perhaps due to higher mortality risk, gender differences in behaviors that affect credit risk, or gender differences in access to credit programs, but it could as well be gender differences in preferences.

The extended family is a potential source of resources for families that want to send their children to school, but feel that they cannot afford it on their own. Angelucci, de Giorgi, Rangel, and Rasul (2010) investigate heterogeneity in *PROGRESA* treatment effects based on the proximity and structure of the recipient's extended family network. They find that households that are connected through intra- or intergenerational links to other households in the same village experience an 9.3 percentage point increase in secondary school enrollment over a 65% baseline enrollment rate, whereas households without such links have treatment effects that are statistically insignificant and close to zero. The proposed mechanism through which the presence of an extended family affects the treatment effect magnitude is the eligibility status of family network members. Since primary school enrollment rates are approximately 90% and there is no average treatment effect on primary school children, the primary school transfers are essentially unconditional. Secondary school enrollment rates increase only among those households with extended family members that are eligible for *PROGRESA* transfers at the primary school level, suggesting that these households draw on extended family resources to offset the costs of sending children to secondary school.

As mentioned above, the Baird, McIntosh, and Özler (2011) study includes an evaluation of an unconditional cash transfer (UCT) program targeting adolescent girls in Malawi, which has a very low enrollment in secondary school (24 percent). A UCT program measures a pure income effect since the cash payments to households are not conditioned on any behavior. The UCT program significantly increased enrollment (by 3.9 pp), but the increase was not as large as in the CCT intervention (8.9 pp).

Edmonds and Schady (2011) evaluate the effects of *Bono de Desarrollo Humano* (BDH), a UCT program in Ecuador, on child labor and school enrollment in that country. The program was targeted to households in the poorest two quintiles of the population

with at least one child age 6-17. In four of Ecuador's 24 provinces the program contained a randomized component, with 1,488 households assigned to receive BDH transfers by lottery.²⁴ The transfer was \$15 per month, equivalent to about 7 percent of average monthly expenditures. The transfer was greater than the increase in the cost of schooling as a student moves from primary school to secondary school, but less than 20 percent of median child labor earnings in the labor market.

Confining their analysis to children aged 11-16 at baseline, the authors find that receiving BDH transfers increased school enrollment by 19 percentage points relative to the counterfactual enrollment rate of 49.1 percent. The transfers reduced paid employment of children by 9.9 percentage points, from 24.3 percent employment among the control group. Participation in unpaid economic activity also fell by 18.7 percentage points relative to the counterfactual participation rate of 55.1 percent. The increase in school enrollment and decline in paid employment were entirely concentrated among children who were students at baseline, and these shifts were somewhat larger among the poorest households in the sample. There was no significant change in hours worked in paid employment conditional on working, potentially because the number of hours worked in the population clusters heavily around 40 hours per week.

Discussion

These studies generally find that income transfers increase students' enrollment and attendance, though in some cases these positive effects are confined to certain sub-groups of children. What is somewhat surprising about the UCT results is their magnitude – they seem too large relative to the income effects on education usually seen. In the Baird, McIntosh and Özler study a 10% increase in household earnings raises secondary school enrollment by more than 16%. The effects in the Edmonds and Schady study are even larger: A 7% increase in household earnings raised secondary school enrollment by 40%. If one takes the implied elasticity seriously, one would expect children in families that live on between 6 and 10 dollars a day per capita at PPP in Ecuador to have secondary school enrollment rates that are 18 times as high as those for families living under 2 dollars a day (they are at least three times as rich and the implied elasticity is six), but in fact the difference is less than a factor of two.²⁵ Even the estimates in Malawi seem much larger than the cross-sectional income elasticities that are observed for other Sub-Saharan African countries (in Tanzania, those under a dollar a day have 40% lower secondary school enrollment rates compared to those whose income is between 2 and 4 dollars a day).²⁶

This raises the issue of whether the potential beneficiaries believe that the UCT is really conditional on school attendance. Both studies try to get at this issue and conclude

²⁴ In contrast to many evaluations of CCT and UCT programs, including Baird, McIntosh, and Özler (2011), the household-level assignment allows observation of lottery winners and losers in the same local area.

²⁵ See "18 country data set," *Poor Economics*. <http://pooreconomics.com/data/country/home>

²⁶ *Ibid.*

that misperceptions of the rules of the UCT are not driving school attendance.²⁷ Perhaps the most obvious alternative is that the families see the UCT as a signal from the government for what they ought to do. This theory would also help explain why, as discussed above, CCTs have very large effects with relatively small monetary amounts but do not have much larger impacts when the incentive amount is increased. Of course it could also be that people have mental accounts and set the money they are getting from this program aside for their children's education. In that case the concavity in the effect of the CCT might come from the fact that there is a limit to how much a family can spend, especially within the range of options that these households have in mind.

Open Questions

1. Why do UCTs have effects much larger than estimates of the income elasticity of the demand for education predict? Is it because parents interpret the program as the government recommendation for how they should behave?
2. Which groups are most affected by UCTs?

E. Information and peer influences

The discussion thus far has implicitly assumed that individuals who obtain higher levels of education will obtain a higher income in the labor market. This subsection first examines the evidence on whether this assumption is supported by the evidence, and then reviews recent studies that examine whether families are always aware of the true costs and benefits of education (and therefore whether, for example, providing information on the returns to education has influence on educational outcomes).

Duflo (2001) exploits a massive school construction campaign in Indonesia from 1973 to 1978 to estimate the returns to schooling for a large cross-section of men born between 1950 and 1972. She links data on adult education and wages to data on the number of schools built in the region of each individual's birth. This intensity of school construction, combined with year of birth indicator variables, serve as instruments for years of schooling in the second-stage wage equation. Duflo estimates that each primary school constructed per 1,000 students led to an average increase in educational attainment of 0.12 to 0.19 years and an increase in adult wages of 1.5-2.7 percent. These estimates imply returns to education—i.e. the percentage increase in earnings per additional year of schooling—of 6.8 to 10.6 percent. This range is similar to most estimates of returns to schooling in developed countries, but somewhat lower than typical estimates for developing countries. Duflo also estimates that the program induced about 1.5 percent of the sample to complete junior secondary school (ending in 9th grade). The program did not affect the educational attainment of individuals who were already completing nine or more years of education.

²⁷ Yet it is true that lottery winners in Ecuador who thought the BDH required schooling were significantly more likely to be enrolled in school.

The study by Spohr (2003), discussed above, also estimates returns to schooling using a policy instrument, in this case the elimination of school fees through the end of junior secondary school in 1968 in Taiwan. Among those reporting positive income from work, instrumental variable estimates of returns to schooling were 5.8% per year for males and 16.4% for females.

The policy question is whether the perceived returns to education are the same as the actual returns. If there is a gap between the two for some set of households, the provision of information about returns can play an important role in determining education outcomes. This subsection examines whether households respond to changes in the (perceived) returns to education, and then discusses a few recent papers that attempt to measure the impact on educational outcomes of providing information on the costs and benefits of education.

Are households more likely to enroll children in school when the returns to schooling are high? More specifically, do they increase children's enrollment in education when those returns increase, or when they are informed that the returns are higher than they assumed? One of the first papers to examine whether households respond to changes in the returns to education was Foster and Rosenzweig (1996). They present evidence that the introduction of high-yielding varieties of rice and wheat in India in the 1960s and 1970s led to increases in the return to education in rural areas of India that were most suited to the adoption of these varieties. More educated farmers may have experienced higher returns to the new varieties because they were better able to apply fertilizer or other inputs in an optimal manner, and/or schooling may have directly facilitated adoption of the new technology. They obtain a district-level estimate of the suitability of rural areas to the adoption of high-yielding varieties, and they find that the districts that were most suited for these varieties, and thus had the largest increases in the returns to education, also had the highest increases in school enrollment.

In a much more recent paper, Jensen (2012) uses experimental estimates from an expansion of female labor market opportunities to estimate the sensitivity of education decisions to actual changes in the returns to education. In particular, Jensen exploits the rapid expansion of the business processing outsourcing industry (BPO) throughout India, largely due to declining costs of providing these services remotely and to Indian regulatory changes.²⁸ These jobs typically have high educational requirements, are well-paid, and prefer female employees. In this study, Jensen provides three years of BPO recruiting services to girls aged 15 to 21 in randomly selected rural villages in India, with the intention of increasing awareness of the newly growing industry in a population in which there was little prior familiarity. In response to the treatment, three years later, women were 2.8 percentage points more likely to have enrolled in vocational or training institutes that offer courses, programs or certification in a range of subjects that include computer programming and English, relative to a baseline enrollment of 0.5 percent. In addition, Jensen tests for increased investment in education in response to anticipated

²⁸ Business processing outsourcing (BPO) industry includes call centers, data entry and management, claims processing, secretarial services, transcription and online technical support, accounting and software development.

higher market returns by examining the schooling choices of younger girls. The results indicated that girls 6 to 17 years of age were 5.0 percentage points more likely to be enrolled in school in treatment villages, compared to a mean baseline enrollment of 73%.

In addition to the impact of information, there is some evidence that access to role models changes beliefs and expectations about what it is possible for women to achieve, in turn increasing the demand for education. Beaman, Duflo, Pande, and Topalova (2012) exploit a 1993 constitutional amendment in India that required that one third of village chief counselor (*pradhan*) positions be reserved at random for a woman in each election cycle. The authors surveyed households in 495 randomly selected villages in Birbhum district, a poor rural district in the state of West Bengal. By the time of data collection, villages' chief counselor positions could have been reserved for women twice (in 1998 and 2003), once (in 1998 or 2003), or never, creating the opportunity to study different levels of exposure to women leaders. They find that adolescent boys in never-reserved villages were 6 percentage points more likely to attend school and 4 percentage points more likely to be able to read and write than their female counterparts, but in the twice-reserved villages the gender gap was not statistically different from zero. Adolescent girls in twice-reserved villages were also 19 percentage points more likely to want to wait until after age 18 to marry, and 8.6 percentage points more likely to want a job that requires an education.

While households may well respond to changes in the returns to education, they may not always have an accurate understanding of what the returns are. Jensen (2010) conducts an experiment on eighth-grade students in the Dominican Republic to investigate the accuracy of perceived returns to education, how individuals update their beliefs when supplied with careful estimates of the returns to education, and the link between the perceived returns to education and schooling decisions.²⁹ In an OLS regression, individual perceived returns to education maintain predictive power over subsequent schooling decisions, even when controlling for academic ability and parental educational attainment, suggesting informational value in reported perceptions. At a randomly selected subset of schools in the study, students had individual meetings with surveyors, who verbally provided information on average earnings by educational level that had been collected through an earlier household survey. Jensen finds that students do update their beliefs regarding the returns to secondary schooling when provided with information, and their updated beliefs translate into greater educational attainment for the treated students. The treatment caused, on average, an increase of 0.20 years of schooling over the next four years in which the students were followed.

Attanasio and Kaufmann (2009) provide additional evidence of the role of individual perceptions regarding the returns to schooling in education decisions. Using

²⁹ An important issue in research of this type is the accuracy of estimates of the returns to education. Indeed, economists often disagree not only on the magnitudes of the returns to education but also on the best method to estimate them. Thus any researchers who attempt research of this type must explain to the students and their parents that economists' estimates, even when carefully done, may not be accurate estimates of what will happen to a particular student. Jensen was very careful to explain this, and other researchers should follow his example.

data from Mexico on individuals' aged 15-25 subjective expectations for future earnings under a number of scenarios for education investment, Attanasio and Kaufmann model the choice to attend high school and college. They find that both youths' and their mothers' expectations matter for high school enrollment, and youths' expectations matter for college enrollment, even with a set of controls for individual and family background.

Low rates of secondary school completion, paired with a substantial earnings premium associated with secondary school completion, provide suggestive evidence that students and their parents may be misinformed regarding the true returns to various education levels. For example, in the Dominican Republic, although earnings of secondary school graduates are over 40% higher than those of primary school graduates, primary school completion is around 80-90% while secondary school completion rates hover around 25-30%. The Jensen (2010) study discussed above uses its experimental design to test whether informational barriers are a key factor determining the low demand for education. From the baseline survey, eighth-grade students, on average, overestimate earnings associated with a primary school education and underestimate earnings associated secondary school and college completion, leading to an underestimate of the returns to education. The intervention, which provided randomly chosen students with information regarding the author-calculated returns to education,³⁰ was effective in decreasing the perceived earnings associated with primary school and increasing perceived earnings associated with secondary school. As discussed above, the treatment caused on average a 0.20 year increase in schooling over the next four years in which the students were followed. It is important to note, however, that the effect is driven by the least poor households: the intervention caused an increase in schooling of 0.33 year for the least poor, while the point estimates are close to zero and statistically insignificant for the poor households. Since both socioeconomic groups updated their perceived returns in similar ways, these results suggest that credit constraints may have prevented the poorest households from increasing their investment in education.

Dinkelman and Martinez (2011) test the impact of providing information on merit-based scholarships and government loans for tertiary education in Chile. They focus on 8th graders and their parents because the majority of Chilean students must choose a high school and field of study by the end of 8th grade. Fifty-six schools were randomly assigned to have their students watch a 15-minute video on the tertiary education experiences of Chileans who grew up in poor families, accessed loans and scholarships for post-secondary education, and went on to professional careers. The video also provided information on specific eligibility cutoffs for these financial resources on Chile's standard high school achievement test. An additional 56 schools were randomly selected to distribute the video to students for home viewing with their parents, and 114 schools served as a comparison group.

Both interventions significantly increased students' knowledge of loan opportunities, and the take-home version significantly increased parents' knowledge. The video also appeared to shift the type of schooling desired by students: those with higher

³⁰ Returns to education calculations are based on national survey data of non-rural households collected by Jensen (2010).

grades were more likely to report that they would study at a college, while those with lower grades were more likely to report that they would study at a vocational training school. The incidence of any absence in the interviewed month among students assigned to either information treatment dropped by 8.8 percentage points relative to the control absenteeism rate of 64%, but six months after the intervention there were no significant effects on test scores or where students enrolled in secondary school.

Loyalka, Song, Wei and Rozelle (2010) use an experimental framework to test the effect of providing information regarding college costs and financial aid on persistence throughout the application process and the likelihood of receiving financial aid. Their sample consists of the best high school in each of the 41 nationally-designated poor counties in the Shaanxi province of China. Students in the science track of randomly selected schools were given a 30-page booklet detailing financial aid programs supported by the government, how to apply to these programs, and additional resources regarding financial aid. The booklet also contained information regarding the costs associated with attending different universities. While there was no statistically significant effect of information provision on the likelihood of receiving need-based grants or the probability of repeating the college entrance exam, students were three percentage points more likely to receive support from a program designed to reduce the transaction and direct costs of college for low-income students.

A related literature on the demand for schooling examines the role of family and peer networks in educational attainment. Family and peer networks can potentially serve as a means both to obtain information regarding educational decisions and to pool financial resources. In addition, social networks generate social norms regarding educational behavior. The papers below exploit the same policy experiment, *PROGRESA*, and generate similar results: educational choices are affected by the decisions of one's peers, particularly close peers.

Bobonis and Finan (2009) test for neighborhood peer effects in secondary school decisions. Part of the implementation of *PROGRESA* included conducting a survey of all households in treatment and control villages in order to construct a welfare index. Only households that were (1) classified as poor according to the welfare index and (2) in treatment villages were eligible to receive cash transfers. Bobonis and Finan compare the secondary school enrollment rates of children classified as non-poor in treatment villages to those of children classified as non-poor in control villages. The results suggest substantial peer effects: they estimate marginally statistically significant spillovers of 5.0 percentage points, relative to a secondary school enrollment rate of 70 percent among non-poor children at the baseline survey. This effect is driven by the non-poor children closest to the welfare index cutoff for eligibility, i.e. the children of relatively poorer households among those that were ineligible for the program. Those children below the median welfare index level among non-poor children experienced an enrollment increase of 5.7 percentage points. Bobonis and Finan also estimate how the rise in the overall village enrollment rate affects the enrollment of non-poor children in treatment villages, instrumenting the village enrollment rate with *PROGRESA* treatment status. The instrumental variable estimates suggest that a 10 percentage point increase in village

enrollment rates increases enrollment of non-poor children by 5 percentage points. Again, the increase in enrollment rates was larger for children of households closer to the welfare index cutoff. There was no evidence of differential spillover effects between girls and boys.

In a closely related paper, Lalive and Cattaneo (2009) utilize the *PROGRESA* experiment to estimate the role of social interactions in schooling decisions. Specifically, they exploit variation in the fraction of a child's classroom peers who are program-eligible in order to estimate how the changing composition of an ineligible student's peers affect his or her school attendance. Confirming the results of Bobonis and Finan (2009), they find that ineligible children attend school more as a result of *PROGRESA*. Lalive and Cattaneo furthermore decompose the effect of *PROGRESA* on eligible children into a direct effect, due to the cash transfers, and an indirect effect, due to social interactions. The social interaction effect is shown to be nearly as important as the cash incentive.

Discussion

As expected, there is credible evidence that additional years of education increase earnings, and households are responsive to changes in the returns to education. It is less clear, however, that students and their parents have accurate information on the economic benefits associated with various levels of education. Indeed, the studies discussed above have found behavioral changes associated with information provision campaigns, leading to the conclusion that informational constraints are an important barrier in educational decision-making.

The provision of information seems to be an inexpensive way to improve the efficiency of educational investment. There are, however, many interesting open questions about what sort of information is most effective.

Open Questions

1. How does the content of the information provided affect school enrollment, performance and attainment? Is it sufficient to give the average rate of return to education or should there be information about the shape of the distribution of returns?
2. Are economists' estimates of the rates of return to additional years of education reasonably accurate? How accurate must they be in order to be helpful, as opposed to potentially harmful, to students and their families?
3. How much emphasis should be on job market information and how much on the returns to education in terms of increased wages later on?
4. Can households distinguish between different quality levels of education, and is there a role for information in helping households make that distinction? Relatedly, are households willing to pay differently for different qualities or types of education?

5. How does the medium of communication for the information affect school enrollment, performance and attainment?
6. How accurate are households' current beliefs regarding the returns to educations? Are their beliefs shaped by local observations, experience or national trends?
7. What are the relative effects of information versus exposure to role models? Are these complements or substitutes?

F. Health

There is scant research on how health conditions affect school enrollment and performance at the post-primary level. The intersection of health and education has been explored much more extensively at the primary level. Simple, low-cost health interventions including school-based deworming, iron fortification, and iodine supplementation have been demonstrated to be highly cost-effective at increasing school participation (Kremer and Holla 2009). A long-run study on deworming in a high-worm-load area of Kenya found that regular deworming treatment increased girls' probability of passing the primary school leavers' exam by nearly 10 percentage points. Additionally, a decade after primary school, deworming increased adult wages among both male and female wage earners (Baird et al. 2012). School meals have also been shown to increase school participation and learning, but through the channel of higher school attendance and achievement on curriculum tests rather than nutritional impacts on student alertness and brain development (Vermeersch and Kremer 2005).

The one study that exists of a health intervention at the post-primary level addresses how female-specific sanitary needs due to menstruation might affect school absenteeism. Oster and Thornton (2011) investigate the hypothesis that adolescent girls may miss school during their menstrual periods if they lack access to modern sanitary products. They provided reusable menstrual cups and training on how to use them to twenty-five randomly selected seventh- and eighth-grade girls at each of four schools in Chitwan, Nepal. The study collected detailed data on the timing of the girls' menstruation, their usage of the cups, and their school attendance. Oster and Thornton find that the girls actually missed very little school due to menstruation: only about half a school day per year, on average. Thus it is not surprising that there was no significant effect on attendance, even though 60 percent of the girls in the treatment group chose to use the cups.

Open Questions

1. To what extent does student health affect educational attainment and performance at the post-primary level?

2. Are nutrition-based interventions still effective—and cost-effective—at the post-primary level, when sensitive periods of development may have passed?
3. Does information or incentives regarding healthy lifestyles and risky behaviors affect educational outcomes?

VI. Existing Research: Supply of Education

It is useful to group the supply side interventions into four categories: incentives and organization, school inputs, and school and teacher quality, selection rules, and pedagogy. We will discuss these one by one.

A. Incentives and Organization

The delivery models for post-primary education tend to be much more diverse than in the case of primary schools. Private, religious, and NGO-run schools and colleges all play important roles; within private schools some are largely government funded while others are either funded by the fees they collect or rely on charitable donations; there are also tutorial centers and distance learning programs, especially at the higher secondary and college levels. Incentives for both teachers and students vary substantially.

One obvious dimension of comparison is between private and public schools. In many developing countries, public colleges are seen as more prestigious than their private counterparts, and their graduates have better educational and labor market outcomes. It is *a priori* ambiguous whether public or private provision of higher education is preferable. Private colleges have an incentive to provide higher quality education in order to compete with public institutions, but at the same time, they have incentives to reduce costs that may undermine quality (Sekhri and Rubinstein 2011).

Sekhri and Rubinstein (2011) find evidence that it is the sorting of better students into public colleges, rather than better value added, that drives higher exit exam scores for public over private college graduates in India. Their identification strategy exploits two features of India's highly regulated tertiary education system. First, admissions to public colleges are based on students' scores on the Senior Secondary School examinations. To be admitted, students must score above a cutoff that varies by field, gender, and caste, and the admissions cutoffs are unknown to the public. The sharp cutoff creates the opportunity to implement a regression discontinuity design. Second, public and private colleges are required to affiliate with a university, and all students in the same field with the same university affiliation are exposed to the same curriculum and exit exams. Therefore, educational outcomes are comparable across public and private college students with the same field and university affiliation. Their sample is based on five years' worth of admissions records for two public and two private colleges in the same urban area. All of the colleges in the area are single-sex, and the authors select one men's and one women's college of each type.

They find that all public college graduates in their sample scored 0.5 standard deviations higher, on average, on exit exams compared to private college graduates. The public-private gap was about a third larger when the sample was restricted to women's colleges. However, this gap disappears completely once entry scores are accounted for, indicating that the observed differences between public and private schools were due to sorting of higher-achieving students into public colleges, among both women and men.

Angrist et al. (2002) use a randomized natural experiment to examine the impact of Colombia's school voucher program, the Programa de Ampliación de Cobertura de la Educación Secundaria (PACES), on school choice, duration of schooling, and test scores. In cities and towns where demand for PACES vouchers exceeded supply, local governments distributed private school vouchers by lottery to eligible low-income students entering the secondary school cycle. Students had to apply and be accepted to secondary school to be deemed eligible. While the voucher initially corresponded to average tuition (US \$190), it was not adjusted for inflation, and thus only partially subsidized tuition. Winners' vouchers were renewed every year unless they were required to repeat a grade.

The authors surveyed three past PACES applicant cohorts and found significant differences in outcomes between voucher winners and losers, especially for girls. Regarding school choice, three years after the lottery voucher winners were 15-16 percentage points more likely to be attending private school than losers, and the effect was larger for girls than for boys (18.2 versus 12.4 percentage points, respectively, both statistically significant). This indicates that the choice between public and private school was sensitive to the change in the price of private school. Voucher winners' school fee payments were US \$52 higher than voucher losers', suggesting that some winners may have used vouchers to attend higher-priced schools. Winners were 10 percentage points more likely to have completed three years of secondary school, primarily because they were less likely to repeat grades. Improvements in test scores were larger and more significant for girls (0.26 standard deviations, significant at 5%) than boys (0.17 s.d., not significant). Voucher winners were also less likely than losers to get married, cohabit, or work while they were teenagers.

Angrist, Bettinger and Kremer (2006) use the same randomized natural experiment to assess the long-term outcomes of the PACES program seven years after winning a voucher. Using administrative records from Colombia's college entrance examination, the ICFES, the authors gathered data on academic achievement and secondary school graduation status. They use ICFES registration as a proxy for graduation, as 90% of graduating students take the exam. Controlling for demographic factors, voucher winners were 6 percentage points more likely to have graduated from secondary school than voucher losers. There were no clear gender differences in impact on graduation, though the baseline graduation rate for boys was lower. Due to selection bias in the simple comparison of test scores of voucher winners and losers, the authors estimate the impact of winning a voucher on ICFES scores using i) Tobit estimates with artificially censored data, and ii) nonparametric bounds for quantile-specific program impacts on the distribution of ICFES scores. After adjusting for selection bias, they find a

positive effect of winning a voucher on ICFES scores, indicating that students who attended private school learned more than those who did not. The impact on exam scores was almost identical for girls and boys.

Private schools differ from public schools in a number of ways, including the incentives faced by teachers and the composition of the peer group. Can we separate these different effects?

Bettinger, Kremer and Saavendra (2010) test whether the PACES vouchers increased educational productivity or simply benefited recipients through peer effects by offering them more desirable peers at others' expense. As voucher applicants had to apply and be accepted to either a private vocational school or a private academic secondary school to be eligible for the lottery, the authors hypothesize that the population of voucher applicants who applied to vocational schools rather than to academic schools likely obtained less desirable peers after winning a voucher. If voucher effects persisted in a context where one group of winners acquired less desirable peers, it would provide evidence that PACES vouchers increased overall educational productivity, rather than just redistributing it. Using a simple linear-in-means model of peer effects, the authors provide evidence that peer quality was significantly lower for voucher winners that applied to vocational schools. Even though they had less desirable peers in terms of observable characteristics, voucher winners in vocational schools were 25% more likely to graduate than voucher losers and had college entrance examination scores one-third of a standard deviation higher than voucher losers. This suggests that the observed effects of PACES vouchers were not simply the result of interaction with better peers and thus offers evidence that the vouchers increased overall educational productivity, which implies that private schools are more effective than public schools. Once again it seems that access to private school, made possible by the vouchers, led to improvements in educational outcomes as measured by graduation rates and exam scores.

Another approach is to look directly at the effect of incentives. Kingdon and Teal (2010) analyze how higher wages for teachers due to teacher union membership affect student proficiency. Utilizing a data set of the exam scores of grade 10 students in 186 Indian English-medium private secondary schools, Kingdon and Teal employ a student fixed effects approach – essentially exploiting exam score variation across subjects for a given student – in order to estimate the effect of a teacher's union membership status on student proficiency. The results suggest that in subjects taught by a unionized teacher, students perform 0.17 to 0.23 standard deviations lower than in subjects taught by a non-unionized teacher, and the effect is negative and significant when estimated separately by student gender. Although they cannot fully rule out that unobservable teacher characteristics correlated with union membership status are driving the results, Kingdon and Teal present some evidence that the magnitude of the selection bias is small. In addition, they estimate a union wage premium of 9.5%. Regarding the mechanisms at work, it is unclear whether the additional compensation associated with union membership is crowding out intrinsic teacher effort, or teacher unobservables correlated with union membership are driving the lower student achievement scores.

Menezes-Filho and Pazello (2007) analyze a Brazilian policy change that increased the share of financial resources states and municipalities must allocate toward teacher wages in order to gauge whether teacher effort is sensitive to wage levels. Teacher wages on average increased due to the policy change, but the extent to which they increased varied by municipality and state since they differed in the initial share of educational resources devoted to teacher wages. The two stage least squares analysis implies that a half standard deviation increase in test scores could be achieved with a doubling of teacher wages. However, again the underlying mechanisms are unclear; it could be that increased pay causes current teachers to work harder (perhaps because they were motivated by the higher pay), or it may be that better teachers are attracted to higher-wage municipalities or states, or perhaps other factors lead to this result.

Another important issue with respect to incentives is whether they should be implemented at the school level or at the level of the individual teacher. Barrera-Osorio and Raju (2011) explore the effects of public subsidies to private schools in Pakistan on school quality, as measured by physical inputs and student-teacher ratios. Due to the vast proliferation of private schools (primary, middle and secondary) in Pakistan in recent years, the growing enrollment of low- and middle-income students in these schools, and the public school system's inability to improve school quality and student outcomes, there is substantial interest in public-private partnerships to improve the provision of education. The program consisted of a monthly per-student cash subsidy. Private schools in Punjab qualified to receive the subsidies by completing a three-step process, including an application, an unannounced school site check, and a student achievement test. Since each school must obtain a minimum pass rate on the student achievement test in order to qualify for the subsidy, the authors are able to compare the outcomes of the schools that barely qualify to those that barely miss the cut off with regression discontinuity techniques. Conditions for maintaining eligibility include, among others, school infrastructure standards and minimum physical input improvements. Using a sharp RD analysis, Barrera-Osorio and Raju find that ten month exposure to the intervention increased the number of teachers by 3.4 (37%), the number of classrooms by 4 (47%), and the number of blackboards by 2.8 (27%) on average for those schools right above the cutoff. Student-teacher ratios were unaffected by the intervention, but this is likely due to enrollment expansion induced by a program eligibility criterion stipulating that schools must eliminate school fees. Though the majority of the schools in the sample were middle schools, the authors do not provide estimates by school level.

Discussion

The studies in this section cover an eclectic range of topics within school organization and incentives. The evidence on public versus private provision is limited and inconclusive. The evidence from India suggests no difference in quality between public and private colleges. However, from the voucher studies in Colombia, we learn that individuals who won vouchers to attend a more expensive private school had better educational outcomes, including higher school graduation rates. The better outcomes, however, might be primarily driven by access to higher quality peers, not by superior provision of education by private schools.

The studies on teacher incentives suggest that these incentives matter for student outcomes, but they do not deal with pay that is tied explicitly to outcomes. In the primary education literature, giving teachers bonus pay based on an objective measure of their attendance has been shown to improve student achievement. Providing bonus pay for teachers based on their students' test scores has had more mixed results; there is evidence of "teaching to the test", so some caution about the design of these incentives is warranted (Kremer and Holla, 2009).

Open questions

1. How does quality differ among private and public post-primary educational institutions? Are these differences driven primarily by organizational differences, profit motives, or student quality?
2. Are voucher programs widely used in developing countries to address inequities in school choice? Do voucher programs in settings other than Colombia improve student outcomes? How do voucher programs affect the full distribution of student outcomes?
3. Is student performance sensitive to teacher pay, without linking teacher pay to student performance?
4. How can teacher incentives be designed at the post-primary level to improve student outcomes? Is a monetary incentive the primary driver of teacher effort, or is teacher recognition more important?
5. Who should provide the teacher incentives? Should the government or an external party like an NGO provide the incentives? Should parents' associations also be involved? Who will measure the outcomes on which the teacher incentives are based (school principal, parents, external observers, or some piece of equipment (cameras, fingerprint reader)?

B. School Inputs and School Quality

One way to get at the issue of school quality is to look at the effect of going to more selective schools. This rests on the assumption that students or their parents can evaluate school quality and therefore the more selective schools are "better". Each of the following four studies makes use of some randomization or quasi-randomization built into the rule by which students get matched with schools.

Pop-Eleches and Urquiola (2011) find evidence that secondary school students in Romania benefit from being placed in higher-level schools and tracks within schools. They apply an RD design that exploits Romania's system for assigning 8th graders to secondary school. Students' high school and academic track (i.e. subject matter) assignments are a deterministic function of their expressed preferences and their "transition score," an average of their performance on a nationwide exam and their grade point average. A centralized process assigns higher-ranking students to their preferred school/track combinations until the available slots fill. Using data on the 2001-2003

admissions cohorts from a nationwide administrative dataset, this admissions process generates about 2,000 quasi-experiments at the school level, and over 6,000 at the track level. The outcome measure is student performance on the Baccalaureate exam, which helps determine whether and where students can attend university.

They find that on average, students who scored just above the cutoff for a more selective school or track score 0.02 to 0.10 standard deviations higher on the Baccalaureate exam than those who just missed the cutoff. In other words, students of similar ability performed better when assigned to better schools, providing evidence that school quality—and not just student aptitude—influenced learning outcomes. However, these effects are quite small, and we cannot tell from this study what aspects of school quality matter most. The authors do not investigate whether there is any heterogeneity in these effects by gender. Using survey data, they find suggestive evidence of behavioral responses to this system, including sorting of teachers with higher certification standards into better-ranked schools and a reduction in parental effort among students who just make the cutoff for a more selective school or track.

Lucas and Mbiti (2011) utilize a similar strategy with data on secondary school admissions in Kenya, but do not find an impact of elite schools. Secondary admissions are based on national primary school exit exam scores and student preferences. The most selective government schools in Kenya are called Nationals Schools, followed by Provincial Schools and District Schools. The three tiers of schools differ dramatically in resources, teacher qualifications, and the breadth of subjects offered. Lucas and Mbiti use a regression discontinuity design to compare students on either side of the National School cutoffs. Because the National schools have district-specific quotas and cutoff scores, the study does not rely only on the lowest-achieving students within each elite school for identification; the marginal admitted student from a higher-achieving school district may be far from the lowest ability student in a National school. The dataset consists of students who took the primary school exit exam in 2004 and the Kenya Certificate of Secondary Education (KCSE) exam in 2008.

Lucas and Mbiti find that on average, graduating from a National School is associated with 0.3 standard deviation higher scores on the KCSE exam. However, the point estimate decreases substantially and becomes statistically insignificant once the sample is limited to more similar students, even at a relatively wide bandwidth around the cutoff. The authors do not find evidence of heterogeneous impacts on students by sex, socioeconomic status, or primary school exit exam scores. They fail to find evidence of differential attrition (i.e. selection into the KCSE exam) around the National School cutoffs, but there does appear to be differential attrition around the Provincial School cutoffs. The majority of students around the National School cutoffs are of quite high ability (2 standard deviations above the mean on initial test scores), so the authors note that the impact of school quality could be quite different on students who are closer to the national average.

Jackson (2010) also exploits discontinuities generated by a rule-based school assignment in Trinidad and Tobago. Once again, secondary school admissions are based

on an algorithm using student preferences and standardized test scores. Jackson obtains test score and school preference data for over 31,000 students in the 2000 cohort of Secondary Entrance Assessment (SEA) takers and matches it to the students' performance in the 2004 and 2005 Caribbean Secondary Education Certification (CSEC) examinations. Jackson's estimates suggest that a student who attends a school where peers have half a standard deviation higher SEA scores will achieve a 0.23 grade point higher score on the CSEC exam (roughly 0.12 standard deviations, or one quarter of the difference between an A and a B). The point estimates are generally larger for female test takers than for males, though the first-stage F statistics for these subsamples suggest that the admissions rule may be a weak instrument. The results for mathematics scores are mixed and mostly insignificant, except for the subsample of female CSEC takers.

Zhang (2009) finds no significant impact from magnet schools on standardized exam scores in Wuhan, China. Magnet schools in this context are semi-private middle schools (grades 7-9) that follow the same curriculum as the free public schools, but have more highly qualified teachers and spend more per pupil than public schools. The magnet schools charge the equivalent of about US \$400 per year in tuition, which is the limit set by the city education council and represents about 13 percent of the average annual disposable income for a family of three in the city.

Because the magnet schools are vastly oversubscribed, a lottery is conducted each year to assign admissions. Zhang's sample includes High School Entrance Exam (HSEE) scores for 13,000 students who participated in admissions lotteries in 2002-2004. While lottery winners and losers are balanced on observable characteristics at baseline, about a third of students admitted to the magnet schools typically gain admission through "back door" processes, such as parents knowing a government official who can influence the principal. To obtain an unbiased estimate of magnet school effects, Zhang uses applicants' lottery status as an instrument for attending a magnet school. Zhang finds that magnet schools had no statistically significant impact on student's HSEE scores, and this result holds for both girls and boys. The higher average HSEE scores observed in magnet schools appear to be the product of student selection outside of the lottery. Zhang speculates on why, if indeed there is no quality advantage for magnet schools, parents might still be willing to pay: the magnet schools may offer something that parents value apart from academic success, such as peer quality; or parents may confuse average student achievement with value added.

These studies have the advantage of credible identification strategies but they do not help us get at what makes a good school. On one hand it could be just selectivity—good schools are good because the best students go there. In one extreme case this could just be a result of signaling—all schools have exactly the same value added, but all the best students compete to get into one of them and therefore graduates from that school are favored in subsequent competition. But there could also be peer effects so that the school that attracts the best students ends up better. Or it could be that the more selective schools have better teachers or better infrastructure. How does one separate these different effects?

Glewwe and Jacoby (1994) used cross-sectional data collected in Ghana in 1988-89 to present correlational evidence of the relationship between teacher and school characteristics and years of schooling and test scores (reading, in English, and math) for students in Ghanaian middle schools (grades 7–10). Of the many school and teacher variables that were examined, most estimated relationships were small and not statistically significant. Five variables were significantly correlated with years of schooling: travel time (a reduction of 2 hours was associated with 2.9 years of additional schooling); teacher experience (raising average experience from 2 to 10 years corresponded with an increase in schooling of 1.5 years); providing blackboards and repairing classrooms that leak when it rains (both associated with about 3 years of additional schooling); and adding a school library (1.7 years of schooling). Only three school variables, and no teacher variables, were directly associated with higher math and reading test scores. Combining the direct and indirect (via the increase in years of schooling) effects, the estimated impact of repairing leaking classrooms, which presumably reduced school closings due to rain, was much larger; the overall (direct plus indirect) impact was an increase of 2.0 standard deviations in reading scores and 2.2 in math scores. Blackboards also had large estimated impacts (direct plus indirect), raising reading scores by 1.9 standard deviations and math scores by 1.8. Adding a library led to smaller increases, 0.3 standard deviations for reading and 1.2 for math scores.

These estimated effects are very large, and since they are based on cross-sectional data they could be biased; for example, despite the large number of controls for school and teacher attributes these estimates could be too large. Indeed, a number of randomized studies of school inputs at the primary level find little impact on student learning, with a limited exception of positive effects of textbooks in Kenya for students who were already in the upper quintiles of test takers (Kremer and Holla, 2009). There are a few studies in post-primary education that focus on variations in individual inputs, resulting from a natural experiment or a field trial.

Barrera-Osorio and Linden (2009) evaluate Computers for Education, a national program in Colombia that takes privately donated and refurbished computers, installs them in public schools, and trains teachers to use them in particular subjects, especially Spanish. The authors conducted a randomized experiment with ninety-seven primary, junior secondary, and secondary schools over two years. They tracked over 5,000 students, who were in grades one through nine at baseline. The student survey included a shortened version of the Colombian national exam, and math and Spanish teachers and school principals were also surveyed.

The authors find almost no significant impact on student test scores or other outcomes. Test scores in Spanish were significantly lower among endline 8th graders in treatment schools, and significantly higher among 9th graders, but the authors attribute these differences to random variation, since no other specifications found significant impacts and there was no difference in teacher-reported computer usage at those grade levels. In fact, there was a significant increase in computer usage only in the two lowest grade levels. The point estimates of impact were larger for girls than boys, but were not significant for either gender. There was fairly serious attrition in the study sample (37

percent), but the authors provide some evidence that this is not a source of bias in their results.

The authors attribute the absence of an effect on learning outcomes to a failure to incorporate the computers into the educational process. Although students in the program schools were 30 percentage points more likely to report using a computer at school in the last week, the only subject in which they appeared to use computers was computer science; there were no significant differences in student-reported computer usage in Spanish or mathematics classes. The authors conclude that introducing technology alone, without links to changes in pedagogy, will not improve the teaching and learning process. The result contrasts with several primary-level studies that show improvements from introducing computer-assisted learning (Kremer and Holla, 2009). The evidence suggests that successful applications of computers help impose a correct curriculum and/or enable students to move through material at their own pace, so this possibility deserves further investigation at the post-primary level.

Greater access to ICT could arguably increase school enrollment by making schools more attractive to students. Cristia, Czerwonko, and Garofalo (2010) use two different empirical strategies to test this hypothesis, estimating the impact of ICT on drop-out and repetition rates at public secondary schools in Peru. First, they exploit a plausibly exogenous increase in the number of computers per student due to a 2004 Inter-American Development Bank (IDB)-funded program that distributed 10 computers to 350 schools. They apply propensity-score reweighting to construct a comparison group, which is similar on observable dimensions but exhibited a relatively flat trend in computer availability per student. They find no effects of computer access on drop-out or repetition. However, the estimates are not very precise, so they also conduct a fixed-effects analysis on a large longitudinal dataset of public urban secondary schools. The dataset tracks a variety of school inputs and student outcomes in 7,319 schools from 2001 through 2007. While this methodology yields much more precise estimates, the authors acknowledge that it may be more susceptible to bias than the results from the policy experiment. In any event, they once again find no significant effects on repetition or drop-out.

Bellei (2009) uses difference-in-differences to estimate the impact of lengthening the school day of secondary school students on the academic performance of 10th grade students in Chile.³¹ In 1996, the Chilean government initiated a switch from a regime of

³¹ Difference-in-differences (DD) is another econometric technique that, like instrumental variables and regression discontinuity designs (see footnotes 17 and 18), attempts to establish a plausibly causal relationship between two variables in the absence of a controlled experiment. In a DD design, researchers compare changes in the outcome over time in the “treatment” group (schools with longer days) to changes over the same period in the untreated group (schools with traditional-length days). This helps control for the fact that test scores likely increase over time even without longer school days. If test scores grow faster (slower) in treated schools than in untreated schools, it suggests that the treatment has a positive (negative) causal effect on test scores. Importantly, interpreting the difference causally depends on the treated and untreated schools being on “parallel trends” in the absence of the treatment. In other words, we are relying on the assumption that there are no other factors, such as socioeconomic differences or other policy interventions, that would cause a systematic difference in the path of test scores across the two groups.

morning and afternoon “shifts” to an extended single session that typically included all of the former morning and half of the former afternoon session. Since this change required a large increase in school infrastructure, the government created a public fund for that purpose and non-randomly phased schools into the program (based on factors such as equitable geographic distribution, the technical viability of each school’s expansion, and student socioeconomic characteristics). Individual schools decided how to allocate the additional time between classroom instruction and extracurricular activities, and a process evaluation indicated that the average participating high school allocated 42 percent of the additional time to academic instruction. The difference-in-differences estimates draw on scores from the 2001 and 2003 national exams administered to 10th graders. A set of 112 high schools that entered the program in 2002 are compared with a comparison group of 647 high schools that by 2003 had still not entered the program.

To investigate the possibility of bias from differential trends between treatment and control groups, Bellei computes sample difference-in-differences in student-level and school-level covariates, and finds no evidence of systematic trend differences between program and control schools. He also fails to find a “program effect” existing before the introduction of the treatment using only the 2001 test data. As further robustness checks, Bellei constructs two additional control groups: one consisting of high schools that entered the program before 2000, and hence were always treated, and the other consisting of schools that experienced half the duration of treatment as the program high schools. Finally, he controls for pre-existing trends by subtracting pre-program difference-in-differences (2001 minus 1998) from the pre/post program difference-in-differences (2003 minus 2001).

Bellei finds small (0.05-0.07 standard deviation) but robust impacts of extended instructional time on language scores. These results are not sensitive to the inclusion of student- and school-level covariates (including the student’s gender), choice of control group, or de-trending. The mathematics test score effects are comparable in magnitude, but are not robust to choice of control group or de-trending. In particular, the impact disappears when the program schools are compared to the “always-treated” control group. The program high schools also experienced a faster pre-program decrease in mathematics scores compared to the control high schools, a finding that “resemble(s) those typically found after the introduction of compensatory programs.” Finally, Bellei finds suggestive evidence that the program effect size was greatest among students in the upper levels of the test score distribution.

Discussion

The research on the effects of attending a more selective school produces mixed results. In some cases, there are small, positive effects of attending a more selective school on students’ subsequent academic performance. In other studies, the effect size is close to zero. These results may in part be a reflection of the methodologies used, since regression discontinuity in effect compares the lowest-performing students at selective schools with the highest-performing students at the less selective schools. However, even in Lucas and Mbiti (2011), which uses a broader range of students for identification due

to district-specific cutoffs in Kenya, there is little apparent effect. Once again, these studies tell us little about what makes for a better school, and there may be benefits from selective schools that are not captured by standardized test scores.

With regard to the determinants of school quality, the message from the few studies on school inputs and quality at the post-primary level appear consistent with what we have learned from the primary school literature. As discussed in Kremer and Holla (2009), input-based interventions generally have either no effect or a small effect, at least when introduced into educational systems with elite-focused curricula and weak teacher incentives. However, inputs that help students move through an appropriate curriculum at their own pace are a promising avenue for further study.

Open Questions

1. Are there specific school or teacher characteristics that affect student performance in a wide variety of settings?
2. Is there a role for ICT in post-primary education? Would tailoring educational technologies (including educational software) to students' specific needs improve academic performance? At the post-primary level, to what extent does training teachers in the use of new ICT affect the impact of introducing that ICT on student performance?
3. Is it possible to adapt the successful para-teacher model to post-primary education? More generally, are there ways to improve and adapt the current models of teacher training, especially in situations where the potential teachers do not have particularly strong academic background?

C. Selection Rules

While most children attend a local primary school near their homes, this is less true for secondary schools and even less true for tertiary education, especially in countries where places in these schools are rationed. Most developing countries have complex rules about who can go to which school. Here we will focus on two examples of such selection—single sex schools, and affirmative action policies for groups that were historically discriminated against.

Proponents of single-sex schools argue that boys and girls learn in different ways, either due to socialization or biological differences such as the timing of neurological development, and that girls may learn better in an environment which they do not have to “compete” with boys. It has also been proposed that single-sex schools may reduce the likelihood of sorting into traditionally male and female subjects. Both of these theoretical arguments suggest efficiency gains from moving more students into single-sex schools. Jackson (2011) notes that existing empirical work on single-sex schooling is likely biased by self-selection and potential confounding differences between single-sex and coeducational schools. As in Jackson (2010), this paper exploits the algorithm used to

assign students to secondary schools in Trinidad and Tobago, which is based on standardized test scores and student preferences. Jackson (2011) uses the rule as an instrument to predict single-sex school attendance. The analysis is confined to public coeducational and single-sex schools that share the same curriculum. However, the single-sex schools in the sample are more selective than the coeducational schools and employ more educated teachers on average, suggesting that the estimated impacts may overstate the pure single-sex schooling effect.

Jackson finds no significant effect, on average, of attending a single-sex school. The basic specification shows a small benefit for attending same-sex schools, but the point estimates become small and insignificant when conditioned on gaining admission to a preferred school of any type. However, single-sex schools appear to benefit students with strong preferences for single-sex schools (i.e. those who list at least three single-sex schools among their top four choices). For those who listed four single-sex schools, attending one of those schools has positive and significant effects on the likelihood of taking the secondary exit exam (a proxy for not dropping out) and the number of exams taken and passed. For most outcomes the positive effects were driven almost entirely by female students. Jackson proposes that the benefits of single-sex schools may be driven by better student-school match rather than an innate advantage of single-sex schools in educational production. Jackson also finds that girls at single-sex schools took fewer math and science courses than girls at co-ed schools, contradicting the notion that single-sex schools reduce gender differences in course selection.

Affirmative action policies for higher education are common throughout the developing world. China, India, South Africa, Brazil and Malaysia have extensively implemented affirmative action policies. Latin American countries generally utilize these policies as part of an effort to expand access to higher education for indigenous groups, while sub-Saharan African countries have implemented them to expand access for women (Lewis and Lockheed 2006). The two papers selected for inclusion in this review deal exclusively with India, though the efficiency and distributional implications may be generalizable to programs in other countries that aim to expand access to disadvantaged groups.

Bertrand, Hanna, and Mullainathan (2010) and Robles and Krishna (2012) exploit India's quota system for historically disadvantaged groups—Scheduled Tribes (ST), Scheduled Castes (SC), and Other Backward Castes (OBC) — to analyze the effects of affirmative action in tertiary education. Under this policy regime, a certain proportion of admissions slots at state-run universities are reserved for each disadvantaged group. Admissions are a deterministic function of one's group and one's score on standardized examinations. Both papers examine issues that are at the heart of vigorous policy debates about affirmative action: *targeting*, the extent to which the policies give access to the truly disadvantaged (as opposed to well-off members of disadvantaged groups); and the *mismatch hypothesis*, the idea that giving disadvantaged students access to a program of study ill-suited to their preparation and credentials may have little impact or actually make them worse off.

Bertrand, Hanna, and Mullainathan (2010) use data from the entrance exams to engineering colleges in one Indian state. To assess targeting, they compare the socio-economic characteristics of applicants who were admitted as a result of the quota policy to those who were displaced, i.e. applicants who were not admitted but would have been in the absence of the policy. To do so, they assume that the size and composition of the applicant pool would have been the same without the policy, and then calculate what the cutoff score would have been, holding the total number of seats at engineering schools constant. They estimate that the average parental income among students admitted thanks to the quotas was only 59 to 72 percent of that of displaced students, suggesting that the quota policy successfully targeted lower-income students. There is also evidence that the quota policies reduced gender diversity: assuming a 50% yield of admitted applicants, 27% of those displaced are female, while only 16% of the students who displaced them are female. This may reflect the greater gender inequality in educational attainment in India's lower castes.

To test for mismatch, these authors estimate the labor market returns to engineering college for both upper- and lower-caste groups. They estimate both OLS and IV regressions, using an indicator variable for "above cutoff for admissions" as an instrument in the latter. The rationale for the IV estimates is potential selection bias based on, among other issues, the presence of liquidity constraints and variation in outside options across caste categories. They find consistently positive returns for both groups, but smaller returns for the lower-caste graduates. While the finding of positive returns for lower caste students is generally contrary to the mismatch hypothesis, the smaller magnitude of the returns for lower caste students does suggest that ignoring the distributional consequences of the policy (which may be important and desirable), the policy leads to aggregate economic losses, since returns for the classes subject to the quota are lower than those for the rest of the population.

Robles and Krishna (2012) use data on the 2008 graduating class at the Indian Institute of Technology (IIT) in Delhi, which has strict quotas for ST and SC students both in general admissions and in academic majors. They investigate the targeting and mismatch issues and evaluate whether disadvantaged students admitted under affirmative action catch up or fall further behind. Like Bertrand and coauthors, they find evidence of successful targeting: the minority students admitted to IIT came from poorer districts than the students they were displacing. However, Robles and Krishna find evidence of mismatch. They compare the wages at graduation of students in selective majors with those of their same-group counterparts in non-selective majors. To account for potential selection into selective majors based on unobserved characteristics that could also affect wages, they jointly estimate two equations relating (1) wages at graduation to the probability of being enrolled in a selective major, and (2) propensity scores for selective majors to observable characteristics and other controls. Within the SC/ST group, they find a significant, negative coefficient on selective majors in the wage regression. In other words, minority students who enroll in selective majors as a result of affirmative action policies appear to earn lower wages than they would have earned if they had chosen a less selective major. Robles and Krishna also find evidence that minority

students not only failed to catch up at IIT, but they fell behind their peers in the general student population, particularly in the selective majors.

Discussion

There is little evidence that single-sex education affects students' academic performance. Studies on selection mechanisms intended to help historically disadvantaged populations are based exclusively on India. From these studies, it is clear that the affirmative action policies are successful at targeting disadvantaged students, but it is unclear whether these students unequivocally benefit from the placement and whether the students displaced by these policies would have benefited more from the placement.

Open Questions

1. What are good selection rules? How much weight should be put on actual preparation for that stage as against innate potential, given that some of the more advantaged children will be better prepared than more talented children from less advantaged backgrounds? How do we separate innate talent from preparation?
2. How should selective policies optimally integrate the beneficiaries into the education system?
3. What methods should be used for selection into vocational versus general education?
4. In countries other than India, what are the targeting and mismatch effects of affirmative action policies?

D. Pedagogy and Curriculum

There is a large body of evidence at the primary level showing that simple interventions can help overcome curricular weaknesses and distortions in teachers' incentives, in particular incentives to focus their attention on guiding the top students through a centrally imposed curriculum while allowing the majority of students to fall behind. As mentioned above, a number of these interventions involved training local volunteers to teach basic literacy and numeracy skills. Hiring local teachers on a short-term contract and splitting classes by initial test scores ("tracking") have also proven effective at improving learning outcomes for both higher- and lower-achieving students. In other studies, information and communications technologies (ICT) that help impose an appropriate curriculum, and/or that allow students to move through material at their own pace, have been shown to improve learning (see Kremer and Holla, 2009 for a review).

There is very little rigorous evidence on what types of curricula or pedagogies are most effective at the post-primary level. Given the greater complexity of material taught in secondary or vocational schools and the qualifications required to teach it, it may not be feasible to replicate interventions with local volunteers and para-teachers at the post-

primary level. There is likely to be an even greater role at the post-primary level for ICT platforms such as computers, tablets, and mobile devices. One finding from the primary education literature that is likely to carry over to the post-primary sphere—and reinforced by Barrera-Osorio and Linden's (2009) study described above—is that simply providing technological inputs is not enough to improve learning. A great deal of research will be needed to assess how to use ICT to help students learn at their own pace and acquire knowledge and skills that are relevant to the labor market.

One study by Aker et al. (forthcoming) gets at these questions. This paper assesses the effect of randomly varying instructional content on adult educational outcomes in Niger. Aker et al. find that including instruction on the use of simple mobile phones with voice and SMS capacity into a basic adult education curriculum substantially improved learning outcomes. Adults in the villages where mobile phone education was implemented scored between 0.19 – 0.26 standard deviations higher on math and writing tests than those in standard adult education classes. The authors attribute this effect to increased student effort and motivation due to the growing importance/relevance of mobile phone technology as a means to communicate inexpensively with one's social network, participate in mobile money programs, and obtain price and labor market information. Students could also practice the basic literacy and numeracy skills required to operate a simple mobile phone outside of the classroom, which may have contributed to improved learning outcomes. The authors point out that mobile technology is more likely to serve as a complement to rather than a substitute for highly educated teachers, though more research is needed to understand the precise mechanisms through which information communication technology increases learning

Blattman, Fiala and Martinez (2011) use an experiment in Uganda to investigate the effects of vocational training. The program they evaluate gave cash grants to young men and women ages 16-35 for vocational school training fees and startup costs for businesses. Applicants for these grants were asked to form a group of 20 to 30 young adults and submit a proposal detailing the vocational training and materials necessary to start their desired business. Among a screened set of applicants, groups were randomly selected for the program and received a large grant, equivalent to \$374 per person. Groups were not subject to rigorous monitoring of how the grant was spent. From self-reported data, however, the majority of the cash grants were utilized for the proposed purpose: two years after the intervention, 79 percent of those who received the treatment had enrolled in vocational training, compared to 17 percent of those in the control, which constitutes an increase of 405 hours of vocational training. There were also positive results for the acquisition of business assets. In terms of economic outcomes, there was a substantial increase in skilled employment, income and consumption. Both males and females increased their employment outside the home, but females more so (50 percent versus 25 percent for males). Other economic impacts are similar among the two genders. Though this treatment was not limited to education, the results provide some evidence that vocational training – when paired with available funds for capital acquisition – has the potential to substantially improve economic outcomes. The question remains whether vocational training alone is equally effective.

Discussion

Research on pedagogy and curriculum is too limited at the post-primary level to draw any general conclusions, and thus much work is needed in this area. The one study that attempts to experimentally vary curriculum components, Aker et al. (forthcoming), provides encouraging results that the direct relevance of the curriculum can improve student outcomes. Similarly, research on vocational education is very scarce, though there is some emerging evidence that vocational education can have positive economic impacts.

Open Questions

1. What is the optimal mix of theoretical knowledge and practical knowledge in post-primary teaching: How much of it should it be computer science rather than computer programming? Spoken and written English rather than English literature? What about soft skills versus hard skills?
2. Do teaching styles matter? How do learning outcomes vary for dynamic, inquiry-based and student-centered learning approaches compared to a rote learning teaching style?
3. What is the role of a tutoring and supplementary teaching, especially given that there are many students who will enter the post-primary stage without having mastered the primary-level content? Should content be taught in the mother tongue or in some international language? And do the answers to these questions depend on what has been done about the demand side and incentives of teachers and students?
4. How to deliver vocational training? How much of it should be hands on practical training and how of it can be classroom instruction? Can virtual models be used for the training? How do you incentivize the teachers?

VII. Summary and Directions for Future Policy-relevant Research

There is a reasonable amount of good research on post-primary education, but the results with the clearest policy implications are concentrated on the demand side. It is clear, for example, that the price of education matters, with the poor being more likely to be priced out of the market. Moreover, there are clear income effects at least in part due to the credit constraints. Yet even these results have limitations; we do not yet understand the magnitudes of either the price effects or the income effects. Some CCTs and UCTs seem to be able to generate much larger effects with small amounts of money than other estimates of price and income effects would suggest. It is true that the CCTs have an incentive component that a straight price reduction does not, but the fact that the amount of money given in the CCT does not seem to have a big effect suggests that there is something more than incentives going on. Figuring out what actually is happening here is very important since it lays out the possibility of designing CCTs that get large effects from nominal sums of money. More generally, it is important to rethink the design of the CCTs. Should it be

one prize/scholarship at the end of the year, or multiple short-term prizes? Does the structure of the prize affect the stronger and weaker students differentially? How about using some of the money as an incentive to matriculate or go to college, or forcing students to save a part of it which could help them finance their future education, as in the insightful paper by Barrera-Osorio, Bertrand, Linden, Perez-Calle (2011)? Should the money go to the family or be given in a way that only the student can have access to it? There is also a more fundamental question: Can CCTs work on a large scale in government systems (as in India) where there is no public exam until 12th grade and attendance statistics are routinely manipulated? It may be that in systems where monitoring is weak, UCTs are more effective.

Also on the demand side there are promising interventions that attempt to inform parents and students about their various educational options, and the likely benefits of them. Clearly more work in this area is warranted since it promises large benefits from very inexpensive interventions. In particular how much emphasis should be on job market information and how much on the returns to education in terms of increased wages later on, or informing parents about what is happening within the school, which might help them support their children more effectively?

On the supply side the evidence base is generally much weaker. We know that more selective schools are not necessarily better, at least for the students who are at the margin of qualification. However, that tells us very little about what makes a good school. On the public/private dimension, the Colombia voucher experiment was quite successful (students who went to private schools did better in life), but this was in a context where there was a supply of pre-existing relatively high quality private schools with some excess capacity. Would that also be true in a place like Pakistan, where the supply of high-quality private schools is extremely limited, especially at the post-primary level, and there is probably not much excess capacity in the system? An interesting laboratory in which to study this would be India's implementation of the Right to Education Law, which requires every private school to admit 25% students from disadvantaged backgrounds who would not otherwise get in. How these underprivileged students will be selected and once selected, become integrated into the schools, given limited capacity, opens up a very large range of exciting potential research questions.

There is also some evidence that teacher and school level incentives matter in the public system, but the evidence is not robust, and this issue can be very politically sensitive. Moreover, very little is known about incentive design. For example, there is some intriguing evidence from primary schools in Andhra Pradesh (India) that the recognition aspect of incentive payments matters as much as the actual cash amount (Muralidharan and Sundararaman 2011). Is this true in the post-primary context, where doing a good job perhaps requires more mastery of the material and relatively less physical energy and enthusiasm? There is also the question of who will provide the incentives: In the Muralidharan and Sundararaman study just mentioned, an NGO was responsible for distributing the incentive payments and teachers claimed that was important because they did not trust the government to honor its commitments. How do incentives get institutionalized? For example, can parents' associations be involved in that process?

There is also some evidence that selection policies play an important role. While affirmative action policies do seem to transfer resources effectively, the design of how to optimally integrate the beneficiaries of these policies into the education system remains an open question, with some evidence that it is not being done optimally. More generally, most developing countries do exercise some selectivity in deciding who is allowed to go on to each further step in the education process beyond primary education, which tends to be offered universally. What are good selection rules? How much weight should be put on actual preparation for that stage as against innate potential, given that some of the more advantaged children will be better prepared than more talented children from less advantaged backgrounds? How do we separate innate talent from preparation?

The evidence is especially thin in the related areas of school inputs. There seem to be only a small number of credible evaluations covering school inputs such as greater use of ICT and extended school hours, and none of them report any great success in raising test scores. On the other hand, there is a great deal of excitement about video-based teaching by expert teachers, say from the Khan Academy, in teaching communities worldwide, and ICT is becoming increasingly sophisticated and cheaper by the day. The ICT-related studies in this review focus on computers; there is very little rigorous evidence on the potential for mobile phones or tablets, which may prove more practical in developing-country settings than computers. At the same time, it is likely that the supply of teachers who can effectively teach science or history at a relatively advanced level is not keeping pace with demand. Given all that, it seems very important to continue to explore the possibilities of using ICT in many alternative modalities, and not to get discouraged by the relatively negative experience thus far. In particular, integration of ICT materials into the overall pedagogy remains a challenge, as does ensuring that the students use the computers or other devices for the designated purpose rather than just playing with them. Other potentially interesting areas to explore within ICT include the use of virtual labs – can they replace a more hands on approach – and the use of e-readers (or for that matter libraries), cell phones, and texting to promote the pleasure of reading and non-curricular learning.

Another very important and, it seems, largely neglected area is teacher training. In particular there is now clear evidence that in primary education lightly trained and minimally paid para-teachers and teacher's aides can promote learning especially among those students who are lagging behind. Is it possible to adapt this successful model in any way in the case of post-primary education? More generally, are there ways to improve and adapt the current models of teacher training, especially in situations where the potential teachers do not have a particularly strong academic background?

Finally, pedagogy is the one area where we could not find almost any credible evidence. The one study on vocational education just shows that there are benefits from vocational training, not that the particular model of vocational training is more effective than other models, or even how it compares to conventional classroom education. Yet there are basic questions about how to deliver vocational training: How much of it should be hands on practical training and how much of it can be classroom instruction? Can virtual models be used for the training? How do you incentivize the teachers? Under the traditional

apprentice system, the trainer typically does not get paid – instead, he or she benefits from the help received from the apprentice, and therefore has an incentive to train the apprentice. The modern vocational school based system is obviously more scalable, but is the training of the same quality?

More generally, little is known regarding effective pedagogies for post-primary education. What is the optimal mix of theoretical knowledge and practical knowledge in post-primary teaching: How much of it should it be computer science rather than computer programming? Spoken and written English rather than English literature? What about soft skills versus hard skills? Do teaching styles matter? How do learning outcomes vary for dynamic, inquiry-based and student-centered learning approaches compared to a rote learning teaching style? What is the role of tutoring and supplementary teaching, especially given that there are many students who will enter the post-primary stage without having mastered the primary-level content? Should content be taught in the mother tongue or in some international language? And do the answers to these questions depend on what has been done about the demand side, including incentives for teachers and students? The answers to these and other related questions have the potential to make huge difference to long run education and employment outcomes, at a time when there seems to be a big push to promote post-primary education.

References

“18 country data set.” *Poor Economics*. Accessed 6 July 2012.

<<http://pooreconomics.com/data/country/home>>

Aker, Jenny, Christopher Ksoll and Travis J. Lybbert. “Can Mobile Phones Improve Learning? Evidence from a Field Experiment in Niger.” Forthcoming, *American Economic Journal: Applied Economics*.

ASER. 2012. *Annual Status of Education Report: Rural*. 16 January.

<http://images2.asercentre.org/aserreports/ASER_2011/aser_2011_report_8.2.12.pdf>

Andrabi, Tahir, Jishnu Das and Asim Ijaz Khwaja. 2011. "Students today, teachers tomorrow ? Identifying constraints on the provision of Education." World Bank Policy Research Working Paper Series 5674.

Angelucci, Manuela, Giacomo de Giorgi, Marcos A. Rangel, and Imran Rasul. 2010. “Family networks and school enrolment: Evidence from a randomized social experiment.” *Journal of Public Economics* 94: 197-221.

Angrist, Joshua, Eric Bettinger, Erik Bloom, Elizabeth King, and Michael Kremer. “Vouchers for Private Schooling in Colombia: Evidence from a Randomized Natural Experiment.” *American Economic Review* 92(5): 1535-1558.

Angrist, Joshua, Eric Bettinger, and Michael Kremer. 2006. “Long Term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia.” *American Economic Review* 96(3): 847-862.

Attansio, Orazio, Emla Fitzsimons, Ana Gomez, Martha Isabel Gutierrez, Costas Meghir, and Alice Mesnard. 2010. "Children's Schooling and Working in the Presence of a Conditional Cash Transfer Program in Rural Colombia." *Economic Development and Cultural Change* 58 (2): 181-210.

Attanasio, Orazio, Costas Meghir, and Ana Santiago. 2011. "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA." *Review of Economic Studies* 79: 37-66.

Attanasio, Orazio P. and Katja M. Kaufmann. 2009. “Educational Choices, Subjective Expectations, and Credit Constraints.” Working Paper, March.

Baez, Javier and Adriana Camacho. 2011. "Assessing the Long-term Effects of Conditional Cash Transfers on Human Capital: Evidence from Colombia." Working Paper (July).

Baird, Sarah, Joan Hamory Hicks, Michael Kremer, and Edward Miguel. 2012. "Worms at Work: Long-Run Impacts of Child Health Gains." Working paper, Harvard University.

Baird, Sarah, Craig McIntosh, and Berk Özler . 2011. "Cash or Condition? Evidence from a Cash Transfer Experiment." *Quarterly Journal of Economics* 126: 1709-1753.

Banerjee, Abhijit. 2004. "Education Policy and the Economics of the Family." *Journal of Development Economics* 74(1): 3-32.

Banerjee, Abhijit, Rukmini Banerji, Esther Duflo, Rachel Glennerster, and Stuti Khemani. 2010. "Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India." *American Economic Journal: Economic Policy* 2(1): 1-30.

Banerjee, Abhijit, Rukmini Banerji, Esther Duflo, and Michael Walton. 2012. "Effective pedagogies and a resistant education system: experimental evidence to improve basic skills in rural India." Working paper.

Banerjee, Abhijit, Shawn Cole, Esther Duflo, and Leigh Linden. 2007. "Remedying Education: Evidence from Two Randomized Experiments in India." *Quarterly Journal of Economics* 122(3): 1235-1264.

Barrera-Osorio, Felipe and Leigh Linden. 2009. "The Use and Misuse of Computers in Education: Evidence from a Randomized Controlled Trial of a Language Arts Program." Working Paper, March.

Barrera-Osorio, Felipe and Dhushyanth Raju. 2011. "Evaluating Public Per-Student Subsidies to Low-Cost Private Schools: Regression Discontinuity Evidence from Pakistan." World Bank Policy Research Working Paper No. 5638 (April).

Barrera-Osorio, Felipe, Leigh L. Linden, and Miguel Urquiola. 2007. "The Effects of User Fee Reductions on Enrollment Evidence from a quasi-experiment." Working Paper, January.

Barrera-Osorio, Felipe, Marianne Bertrand, Leigh Linden, and Francisco Perez-Calle. 2011. "Improving the Design of Conditional Transfer Programs: Evidence from a Randomized Education Experiment in Colombia." *American Economic Journal: Applied Economics* 3(April): 167-195.

Beaman, Lori, Esther Duflo, Rohini Pande, and Petia Topalova. 2012. "Female Leadership Raises Aspirations and Education Attainment for Girls: A Policy Experiment in India." *Science* 335(6068): 582-586.

Becker, Gary. 1964. *Human Capital*. New York: Columbia University Press (for National Bureau of Economic Research).

Behrman, Jere, Susan Parker, and Petra Todd. 2011. "Do Conditional Cash Transfers for Schooling Generate Lasting Benefits? A Five-Year Followup of PROGRESA/Oportunidades." *Journal of Human Resources* 46(1):93-122.

Bellei, Cristian. 2009. "Does lengthening the school day increase students' academic achievement? Results from a natural experiment in Chile." *Economics of Education Review* 28: 629-640.

Bertrand, Marianne, Rema Hanna, and Sendhil Mullanaitan. 2010. "Affirmative action in education: Evidence from engineering college admissions in India." *Journal of Public Economics* 94: 16-29.

Bettinger, Eric, Michael Kremer, and Juan E. Saavedra. 2010. "Are Educational Vouchers Only Redistributive?" *The Economic Journal* 120: F204-F228.

Blattman, Christopher, Nathan Fiala, and Sebastian Martinez. 2011. "Can employment programs reduce poverty and social instability? Experimental evidence from a Ugandan aid program." Working Paper.

Blimpo, Moussa, 2010. "Team Incentives for Education in Developing Countries: A Randomized Field Experiment in Benin." Working Paper, March.

Bobonis, Gustavo J. and Frederico Finan. 2009. "Neighborhood Peer effects in Secondary School Enrollment Decisions." *Review of Economics and Statistics* 91(4): 695-716.

Borkum, Evan. 2012. "Can Eliminating School Fees in Poor Districts Boost Enrollment? Evidence from South Africa." *Economic Development and Cultural Change* 60(2): 359-398.

Center for Universal Education. 2011. *A Global Compact on Learning: Taking Action on Education in Developing Countries*. Washington, DC: Brookings Institution.

Chaudhury, Nazmul and Dilip Parajuli. 2006. "Conditional Cash Transfers and Female Schooling: The Impact of the Female School Stipend Program on Public School Enrollments in Punjab, Pakistan." *The World Bank Impact Evaluation Series* No. 9 (December 2006).

Cristia, Julian P., Alejo Czerwonko, and Pablo Garofalo. 2010. "Does ICT Increase Years of Education? Evidence from Peru." Inter-American Development Bank Working Paper OVE/WP-01/10. May.

Dinkelman, Taryn and Claudia Martinez. 2011. "Investing in Schooling in Chile: The role of information about financial aid for higher education." Working Paper, 24 March.

Duflo, Esther. 2001. "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment." *The American Economic Review* 94(4): 795-813.

Edmonds, Eric V. 2006. "Child labor and schooling responses to anticipated income in South Africa." *Journal of Development Economics* 81: 386-414.

Edmonds, Eric V. and Norbert Schady, 2011. "Poverty Alleviation and Child Labor." Forthcoming in *American Economic Journal: Economic Policy*.

Filmer, Deon and Norbert Schady. 2009a. "School Enrollment, Selection and Test Scores." *World Bank Policy Working Paper* 4998 (July).

Filmer, Deon and Norbert Schady. 2009b. "Are there Diminishing Returns to Transfer Size in Conditional Cash Transfers?" *World Bank Policy Working Paper* 4999 (July).

Foster, Andrew, and Mark Rosenzweig. 1996. "Technical Change and Human Capital Returns and Investments: Evidence from the Green Revolution." *American Economic Review* 86(4): 931-953.

Glewwe, Paul. 1991. "Schooling, Skills, and the Returns to Government Investment in Education: An Exploration Using Data from Ghana." Living Standards Measurement Study, Working Paper No. 76.

Glewwe, Paul and Hanan Jacoby. 1994. "Student Achievement and Schooling Choice in Low income Countries: Evidence from Ghana." *Journal of Human Resources* 29(3): 843-864.

Glewwe, Paul and Ana Lucia Kassouf. 2012. "The Impact of the Bolsa Escola/Familia Conditional Cash Transfer Program on Enrollment, Dropout Rates, and Grade Promotion in Brazil." *Journal of Development Economics* 97: 505-517.

Glewwe, Paul W., Eric A. Hanushek, Sarah D. Humpage, and Renato Ravina, 2013. "School Resources and Educational Outcomes in Developing Countries: A Review of the Literature from 1990 to 2010." Forthcoming in *Education Policy in Developing Countries*. University of Chicago Press.

Gurgand, Marc, Adrien Lorraineau, and Thomas Melonio. 2011. "Student Loans: Liquidity Constraint and Higher Education in South Africa." Agence Française de Développement Working Paper No. 117, September.

Hanushek, Eric, and Lei Zhang. 2009. "Quality-Consistent Estimates of International Schooling and Skill Gradients." *Journal of Human Capital* 3(2): 107-143.

He, Fang, Leigh Linden, and Margaret MacLeod. 2008. "How to Teach English in India: Testing the Relative Productivity of Instruction Methods within the Pratham English Language Education Program" Columbia University: Working Paper.

Innovations for Poverty Action. 2011. "Scaling Up." Accessed 6 July 2012. <
<http://www.poverty-action.org/remedialeducation/scalingup>>

Jackson, Kirabo C. 2010. "Do Students Benefit from Attending Better Schools? Evidence from rule-based student assignments in Trinidad and Tobago." *The Economic Journal*, 120:1399-1429.

Jackson, C. Kirabo. 2011. "Single-Sex Schools, Student Achievement, and Course Selection: Evidence from Rule-Based Student Assignments in Trinidad and Tobago." NBER Working Paper 16817 (February).

Jensen, Robert. 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* 125 (2): 515-548.

Jensen, Robert. 2012. "Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India." *Quarterly Journal of Economics* 127(2): 753-792.

Kaufmann, Katja (2012) "Understanding the Income Gradient in College Attendance in Mexico: The Role of Heterogeneity in Expected Returns." Working Paper.

Kingdon, Geeta and Francis Teal. 2010. "Teacher unions, teacher pay and student performance in India: A pupil fixed effects approach." *Journal of Development Economics* 91: 278-288.

Kremer, Michael and Alaka Holla. 2009. "Improving Education in the Developing World: What Have We Learned from Randomized Evaluations?" In Kenneth J. Arrow and Timothy F. Bresnahan, eds. *Annual Review of Economics*, vol. 1.

Kremer, Michael, Miguel, Edward and Rebecca Thornton. 2009. "Incentives to Learn." *Review of Economics and Statistics*, 91(3): 437-456.

Lalive, Rafael and M. Alejandro Cattaneo. 2009. "Social Interactions and Schooling Decisions." *Review of Economics and Statistics*, 91(3): 457-477.

Lazear, Edward P. 2003. "Teacher incentives." *Swedish Economic Policy Review* 10(3): 179-214.

Lewis, Maureen and Marlaine Lockheed. 2006. *Inexcusable Absence: Why 60 Million Girls Still Aren't In School and What to Do About It*. Washington, DC: Center for Global Development.

Lochner, Lance. 2011. "Non-Production Benefits of Education: Crime, Health, and Good Citizenship." In *Handbook of The Economics of Education*, vol 4., Eric A. Hanushek, Stephen Machin and Ludger Woessmann Eds. pp. 183-282.

Loyalka, Prashant, Yingquan Song, Jianguo Wei and Scott Rozelle, 2010. "Information, College Decisions and Financial Aid: Evidence from a Cluster-Randomized Control Trial in China." Working Paper, November.

Lucas, Adrienne M. and Isaac M. Mbiti. 2011. "Elite Secondary Schools and Student Achievement: Regression Discontinuity Evidence from Kenya." Working Paper.

Menezes-Filho, Naercio and Elaine Pazello. 2007. "Do teachers' wages matter for proficiency? Evidence from a funding reform in Brazil." *Economics of Education Review* 26: 660-672.

Mincer, Jacob A. 1970. "The Distribution of Labor Incomes: A Survey." *Journal of Economic Literature*, March.

Mincer, Jacob A. 1974. *Schooling, Experience, and Earnings*. New York: Columbia University Press.

Mulligan, Casey B. 1999. "Galton versus the Human Capital Approach to Inheritance." *Journal of Political Economy* 107(6) (December): S184-S224.

Muralidharan, Karthik, and Venkatesh Sundararaman. 2011. "Teacher Performance Pay: Experimental Evidence from India." Working Paper, 5 January.

Murnane, Richard J., John B. Willett, Yves Duhaldeborde, and John H. Tyler. 2000. "How important are the cognitive skills of teenagers in predicting subsequent earnings?" *Journal of Policy Analysis and Management* 19(4): 547-568.

Oster, Emily and Rebecca Thornton. 2011. "Menstruation, Sanitary Products, and School Attendance: Evidence from a Randomized Evaluation." *American Economic Journal: Applied Economics* 3(1): 91-100.

Ozier, Owen. 2010. "The Impact of Secondary Schooling in Kenya: A Regression Discontinuity Analysis." Working Paper, October.

Pop-Eleches, Cristian and Miguel Urquiola. 2011. "Going to a Better School: Effects and Behavioral Responses." NBER Working Paper 16886, March.

Robles, Veronica C. Frisancho and Kala Krishna. 2012. "Affirmative Action in Higher Education in India: Targeting, Catch Up, and Mismatch at IIT-Delhi." NBER Working Paper 17727, January.

Schady, Norbert and Maria Caridad Araujo. 2006. "Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador." *The World Bank Impact Evaluation Series* No. 3 (June).

Schultz, Paul. 2002. "Why governments should invest more to educate girls." *World Development* 30: 207-25.

Schultz, Paul. 2004. "School Subsidies for the Poor: Evaluating the Mexican *PROGRESA* Poverty Program." *Journal of Development Economics* 74: 199-250.

Sekhri, Sheetal and Yona Rubenstein. 2011. "Public Private College Educational Gap in Developing Countries: Evidence on Value Added versus Sorting from General Education Sector in India." Working Paper, November.

Solis, Alex. 2011. "Credit Constraints for Higher Education." Working Paper, October 21. (NEUDC 2011)

Spohr, Chris A. 2003. "Formal schooling and workforce participation in a rapidly developing economy: evidence from 'compulsory' junior high school in Taiwan." *Journal of Development Economics* 70(2): 291-327.

Todd, Petra E. and Kenneth I. Wolpin. 2006. "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility." *American Economic Review* 96(5): 1384-1417.

UNESCO, 2011. "The Hidden Crisis: Armed Conflict and Education." Education for All Global Monitoring Report 2011.
<<http://unesdoc.unesco.org/images/0019/001907/190743e.pdf>>

UWEZO, 2011a. *The 2011 Uwezo Kenya National Report*. 18 July. <
<http://www.uwezo.net/index.php?i=215>>

UWEZO, 2011b. *The 2011 Uwezo Tanzania National Report*. 19 September. <
<http://www.uwezo.net/index.php?i=264>>

UWEZO, 2011c. *The 2011 Uwezo Uganda National Report*. 18 August. <
<http://www.uwezo.net/index.php?i=244>>

Vermeersch, Christel, and Michael Kremer. 2005 "School Meals, Educational Achievement and School Competition: Evidence from a Randomized Evaluation." The World Bank, Policy Research Working Paper Series: 2523.

World Bank, 2012a. "World Development Indicators." Data retrieved 6 July 2012. <
<http://data.worldbank.org/data-catalog/world-development-indicators>>

World Bank, 2012b. *World Development Indicators 2012*. Washington, DC: The World Bank.

Wydick, Bruce, Paul Glewwe, and Laine Rutledge. 2013. "Does International Child Sponsorship Work? A Six-Country Study of Impacts on Adult Life Outcomes." *Journal of Political Economy*. Forthcoming.

Zhang, Hongliang. 2009. "Magnet Schools and Student Achievement: Evidence from a Randomized Natural Experiment in China." Working Paper, December.

Table 1 – Primary and Secondary Gross Enrollment Rates: 1980 to 2010

Region	Primary			Secondary		
	1980	1995	2010	1980	1995	2010
East Asia and Pacific	111	115	111	43	65	76
Latin American and Caribbean	106	111	117	42	53	90
Middle East and North Africa	87	97	102	42	64	72
South Asia	76	99	110	27	49	55
Sub-Saharan Africa	78	75	100	14	27	36

Source: World Bank (1998, 2012b)

Table 1 (Continued) – Tertiary Gross Enrollment Rates: 1980 to 2010

Region	Tertiary		
	1980	1995	2010
East Asia and Pacific	X	X	25
Latin American and Caribbean	14	15	37
Middle East and North Africa	11	15	27
South Asia	5	6	11
Sub-Saharan Africa	1	...	6

Source: World Bank (1998, 2012b)

Table 2 – Primary School Completion Rates: 1980 to 2010

Region	1991	2010
East Asia and Pacific	100	97
Latin American and Caribbean	83	102
Middle East and North Africa	77	88
South Asia	76	86
Sub-Saharan Africa	50	67

Source: World Bank (2002, 2012b)

Table 3 –Primary Gross Enrollment Rates by Region and Gender

Region	1999		2009	
	M	F	M	F
East Asia and Pacific	110	110	111	112
Latin American and Caribbean	123	120	119	115
Middle East and North Africa	103	91	106	98
South Asia	97	80	113*	107*
Sub-Saharan Africa	84	71	104	95

Source: World Bank (2012a)

* *indicates data from 2008*

Table 3 (cont) –Secondary Gross Enrollment Rates by Region and Gender

Region	1999		2009	
	M	F	M	F
East Asia and Pacific	61	57	74	78
Latin American and Caribbean	77	83	86	93
Middle East and North Africa	69	61	75	69
South Asia	48	35	58	51
Sub-Saharan Africa	27	22	40	32

Source: World Bank (2012a)

Table 3 (cont) –Tertiary Gross Enrollment Rates by Region and Gender

Region	1999		2009	
	M	F	M	F
East Asia and Pacific	11	9	24	26
Latin American and Caribbean	19	23	33	41
Middle East and North Africa	24	17	28	27
South Asia	9	6	13*	9*
Sub-Saharan Africa	5	3	8	5

Source: World Bank (2012a)

* *indicates data from 2008*

**Table 4 – Scores on Internationally Comparable Tests, 2000 to 2009
(15 year old students)**

Country	Subject	2000	2003	2006	2009
Argentina	Reading	418		374	398
	Mathematics			381	388
Brazil	Reading	396	403	393	412
	Mathematics		356	370	386
Chile	Reading	410		442	449
	Mathematics			411	421
Colombia	Reading			385	413
	Mathematics			470	481
Indonesia	Reading	371	382	393	402
	Mathematics		360	381	371
Jordan	Reading			401	405
	Mathematics			384	387
Mexico	Reading	422	400	410	425
	Mathematics		385	406	419
Peru	Reading	327			370
Thailand	Reading	431	420	417	421
	Mathematics		417	417	419
Tunisia	Reading		375	380	404
	Mathematics		359	365	371
Turkey	Reading		375	380	404
	Mathematics		423	424	445
Uruguay	Reading		434	413	426
	Mathematics		422	427	427

Source: OECD (2000, 2003, 2006, 2009)