

**Human capital in growth regressions:
how much difference does data quality make?
An update and further results**

Angel de la Fuente*

Instituto de Análisis Económico (CSIC)

and

Rafael Doménech

Universidad de Valencia

First version: January 2000

This version: July 2002

Abstract

We construct estimates of educational attainment for a sample of OECD countries using previously unexploited sources. We follow a heuristic approach to obtain plausible time profiles for attainment levels by removing sharp breaks in the data that seem to reflect changes in classification criteria. We then construct indicators of the information content of our series and a number of previously available data sets and examine their performance in several growth specifications. We find a clear positive correlation between data quality and the size and significance of human capital coefficients in growth regressions. Using an extension of the classical errors in variables model, we construct a set of meta-estimates of the coefficient of years of schooling in an aggregate Cobb-Douglas production function. Our results suggest that, after correcting for measurement error bias, the value of this parameter is above 0.50.

Keywords: human capital, growth, measurement error

JEL Classification: O40, I20, O30, C19

* This paper is part of a research project cofinanced by the European Fund for Regional Development and Fundación Caixa Galicia. We gratefully acknowledge additional support from the Spanish Ministry of Science and Technology through CICYT grants SEC99-1189 and SEC99-0820 and from the OECD. We would like to thank Georges Lemaitre and Jean-Luc Heller of the Statistics and Indicators division of the OECD and their correspondents in the national statistical offices of member countries who have contributed greatly to the improvement of this data set. We are also grateful to Jon Temple, Antonio Ciccone, Jong-Wha Lee, José Emilio Bosca, Anna Borkowsky (Swiss Federal Statistical Office) and Gunilla Dahlen (Statistics Sweden) for their helpful comments and suggestions, and to María Jesús Freire and Juan Antonio Duro for their competent research assistance.

Correspondence to: Angel de la Fuente, Instituto de Análisis Económico (CSIC), Campus UAB, 08193, Bellaterra, Barcelona, Spain. Tel: 34-93-580-6612. Fax: 93-580-1452. E-Mail address: Angel.DeLafuente@uab.es. or Rafael Doménech, Departamento de Análisis Económico, Universidad de Valencia, Campus de los Naranjos, 46022 Valencia, Spain; email: Rafael.Domenech@uv.es

1. Introduction

Recent empirical investigations of the contribution of human capital accumulation to economic growth have often produced discouraging results. Educational variables frequently turn out to be insignificant or to have the "wrong" sign in growth regressions, particularly when these are estimated using first-differenced or panel specifications.¹ The accumulation of such negative results in the recent literature has fueled a growing skepticism about the role of schooling in the growth process, and has even led some researchers (notably Pritchett, 1999) to seriously consider possible reasons why educational investment may fail to contribute to productivity growth.

In this paper we provide evidence that these counterintuitive results on human capital and growth can be attributed to deficiencies in the data, and show that improvements in data quality lead to larger and more precise estimates of schooling coefficients in growth regressions. In the first part of the paper, we review the main schooling data sets that have been used in the empirical growth literature and document a number of suspicious features and inconsistencies that suggest that these data contain substantial measurement error. Next, we attempt to increase the signal-to-noise ratio in the data by constructing new schooling series for a sample of 21 OECD countries. These series make use of previously unexploited sources, including unpublished data supplied by the OECD and by a number of national statistical agencies, and rely on a heuristic procedure to obtain plausible time profiles for attainment levels by removing sharp breaks in the data that can only be due to changes in classification criteria. Following Krueger and Lindhal (2001), we use estimates of reliability ratios to measure the amount of measurement error in different cross-country data sets and find that, for the case of the OECD sample we work with, our series have the highest information content of all those available, followed closely by the attainment estimates constructed by Cohen and Soto (2001).

In the final part of the paper, we systematically compare the performance of our series and several other attainment data sets in a number of standard growth specifications and find a clear positive correlation between estimated schooling coefficients and reliability ratios. Finally, we use an extension of the classical errors-in-variables model to construct estimates of reliability ratios that allow for correlated measurement errors and to obtain a set of meta-estimates of the coefficient of schooling that correct for attenuation bias. The exercise suggests that the coefficient of years of schooling in an aggregate Cobb-Douglas production function is above 0.50.

¹ See among others Knight, Loayza and Villanueva (1993), Benhabib and Spiegel (1994), Islam (1995), Caselli et al (1996), Hamilton and Monteagudo (1998) and Pritchett (1999).

2. International data on educational attainment: a brief survey and some worrisome features

This section reviews the available cross-country data on educational attainment levels and documents some of the problems they display. We focus primarily on panel data sets that were available as of the mid 1990s because most of the existing empirical studies on the growth effects of human capital have relied on these series.

Most governments gather information on a number of educational indicators through population censuses and labour force and specialized surveys. Various international organizations collect these data and compile comparative statistics that provide easily accessible and (supposedly) homogeneous information for a large number of countries. Perhaps the most comprehensive regular source of international educational statistics is UNESCO's *Statistical Yearbook*. This publication provides reasonably complete yearly time series on school enrollment rates by level of education for most countries in the world and contains some data on the educational attainment of the adult population, government expenditures on education, teacher/pupil ratios and other variables of interest. Other UNESCO publications contain additional information on educational stocks and flows and some convenient compilations. Other useful sources include the UN's *Demographic Yearbook*, which also reports educational attainment levels by age group, and the IMF's *Government Finance Statistics*, which provides data on public expenditures on education. Finally, the OECD also compiles educational statistics both for its member states and occasionally for larger groups of countries.²

The UNESCO enrollment series have been used in a large number of empirical studies of the link between education and productivity. In many cases this choice reflects the easy availability and broad coverage of these data rather than their theoretical suitability for the purpose of the study. Enrollment rates can probably be considered an acceptable, although imperfect, proxy for the flow of educational investment. On the other hand, these variables are not necessarily good indicators of the existing stock of human capital since average educational attainment (which is often the more interesting variable from a theoretical point of view) responds to investment flows only gradually and with a very considerable lag.

In an attempt to remedy these shortcomings, a number of researchers have constructed data sets that attempt to measure directly the educational stock embodied in the population or labour force of large samples of countries. One of the earliest attempts in this direction is due to Psacharopoulos and Arriagada (P&A, 1986) who, drawing on earlier work by Kaneko (1986), report data on the educational composition of the labour force in 99 countries and provide estimates of average years of schooling. In most cases, however, P&A provide only one observation per country.

More recently, there have been various attempts to construct more complete data sets on educational attainment that provide broader temporal coverage and can therefore be used in growth accounting and other empirical exercises. This requires panel data for as many countries and years as possible.

² See especially its annual *Education at a Glance* report.

2.1. Educational data bases: coverage and construction

The existing panel data sets on educational attainment have been constructed by combining the available data on attainment levels with the UNESCO enrollment figures to obtain series of average years of schooling and of the educational composition of the population or the labour force. Enrollment data are transformed into attainment figures through a perpetual inventory method or some short-cut procedure that attempts to approximate it. We are aware of the following studies:

- **Kyriacou (1991)** provides estimates of the average years of schooling of the labour force (H) for a sample of 111 countries. His data cover the period 1965-1985 at five-year intervals. He uses UNESCO data and P&A's attainment figures to estimate an equation linking H to lagged enrollment rates. This equation is then used to construct an estimate of H for other years and countries.
- **Lau, Jamison and Louat (1991) and Lau, Bhalla and Louat (1991)**. These studies use a perpetual inventory method and annual data on enrollment rates to construct estimates of attainment levels for the working-age population. Their perpetual inventory method uses age-specific survival rates constructed for representative countries in each region but does not seem to correct enrollment rates for dropouts or repeaters. "Early" school enrollment rates are estimates constructed through backward extrapolation of post-1960 figures. They do not use or benchmark against available census figures.
- **Barro and Lee (B&L 1993)** construct attainment series by combining census data and enrollment rates, both taken primarily from UNESCO. To estimate attainment levels in years for which census data are not available, they use a combination of interpolation between available census observations (where possible) and a perpetual inventory method that can be used to estimate changes from nearby (either forward or backward) benchmark observations. Their version of the perpetual inventory method makes use of data on gross enrollments³ and on the age composition of the population (to estimate survival rates). The data set contains observations for 129 countries and covers the period 1960-85 at five-year intervals. Besides the average years of education of the population over 25, Barro and Lee report information on the fraction of the (male and female) population that has reached and completed each educational level. In a more recent paper (B&L, 1996), the same authors present an update and extension of their previous work that has been used in a large number of empirical applications. The revised database, which is constructed following the same procedure as the previous one (except for the use of net rather than gross enrollment rates), extends the attainment series to 1990 and to the population over 15 years of age, and provides new information on quality indicators such as the pupil/teacher ratio, public educational expenditures per student and the length

³ The gross enrollment rate is defined as the ratio between the total number of students enrolled in a given educational level and the size of the population which, according to its age, "should" be enrolled in the course. The net enrollment rate is defined in an analogous manner but counting only those students who belong to the relevant age group. Hence, older students (including repeaters) are excluded in this second case.

of the school year. Some further extensions and refinements of this data base have been made available by the authors in recent years and are discussed in Barro and Lee (2000) and Lee and Barro (2001). The latest update of this data set uses gross enrollment rates corrected for repeaters to fill in missing observations, corrects for observed changes in the duration of educational cycles, and includes some information on student achievement as measured by scores in standardized international tests.

- **Nehru, Swanson and Dubey (NSD 1995)** follow roughly the same procedure as Lau, Jamison and Louat (1991) but introduce several improvements. The first one is that Nehru et al collect a fair amount of enrollment data prior to 1960 and do not therefore need to rely as much on the backward extrapolation of enrollment rates. Secondly, they make some adjustment for grade repetition and drop-outs using the limited information available on these variables.

We can divide these studies into two groups according to whether they make use of both census attainment data and enrollment series or only the latter. The first set of papers (Kyriacou and Barro and Lee) relies on census figures where available and then uses enrollment data to fill in the missing values. Kyriacou uses a simple regression of educational stocks on lagged flows to estimate the unavailable levels of schooling. This procedure is valid only when the relationship between these two variables is stable over time and across countries, which seems unlikely although it may not be a bad rough approximation, particularly within groups of countries with similar population age structures. In principle, Barro and Lee's procedure should be superior to Kyriacou's because it makes use of more information and does not rely on such strong implicit assumptions. In addition, these authors also choose their method for filling in missing observations on the basis of an accuracy test based on a sample of 30 countries for which relatively complete census data are available.

The second group of papers (Louat et al and Nehru et al) uses only enrollment data to construct time series of educational attainment. The version of the perpetual inventory method used in these studies is a bit more sophisticated than the one in the first version of Barro and Lee, particularly in the case of Nehru et al.⁴ On the other hand, these studies completely ignore census data on attainment levels. To justify this decision, Nehru et al observe that census publications typically do not report the actual years of schooling of individuals (only whether or not they have completed a certain level of education and/or whether they have started it) and often provide information only for the population aged 25 and over. As a result, there will be some arbitrariness in estimates of average years of schooling based on these data, and the omission of the younger segments of the population may bias the results, particularly in LDCs, where this age group is typically very large and much more educated than older cohorts. While this is certainly true and may call for some adjustment of the

⁴ Differences across these studies have to do with the correction of enrollment rates for dropouts and repeaters and with the estimation of survival probabilities. Latter versions of Barro and Lee have improved the treatment of the first of these issues.

census figures on the basis of other sources, in our opinion it hardly justifies discarding the only direct information available on the variables of interest.

2.2. A closer look at the data for an OECD sample

Methodological differences across different studies would be of relatively little concern if they all gave us a consistent and reasonable picture of educational attainment levels across countries and of their evolution over time. As we will see presently, this is not the case. One problem is that there are very significant differences across sources in terms of the relative positions of different countries. Although the various studies generally coincide when comparisons are made across broad regions (e.g. the OECD vs. LDCs in various geographical areas), the discrepancies are very important when we focus on the group of industrialized countries. Another cause for concern is that practically all available data on educational stocks and flows, including UNESCO's enrollment series, present anomalies which, to some extent, raise doubts about their accuracy and consistency. In particular, the schooling levels reported for some countries do not seem very plausible, while others display extremely large changes in attainment levels over periods as short as five years (particularly at the secondary and tertiary levels) or extremely suspicious trends.⁵

To illustrate these problems and to get some feeling for the overall reasonableness of the existing data, in this section we will take a closer look at the Barro and Lee (B&L 1996) and Nehru et al (NSD 1995) data sets. These are the most sophisticated data sets within each of the groups of studies identified in the previous section that were available in the mid 1990s, and they have been used in many growth studies. As in the rest of the paper, we will concentrate on a sample of OECD countries. One of the main reasons for this choice is that educational statistics for this set of advanced industrial nations are presumably of decent quality. Any deficiencies we find in them are likely to be compounded in the case of poorer countries.

Table 1: Correlation among alternative estimates of average schooling

	NSD	P&A	B&L	Kyr
<i>Nehru et al (NSD)</i>	1			
<i>Psch. and Arr. (P&A)</i>	0.84	1		
<i>Barro and Lee (B&L 93)</i>	0.81	0.92	1	
<i>Kyriacou (Kyr)</i>	0.89	0.86	0.89	1

- Source: Nehru et al (1995).


The degree of consistency between the various sources varies a lot depending on the level of aggregation we consider. Table 1, taken from NSD (1995), shows that the overall correlation (computed over common observations) between different estimates of average years of schooling is reasonably high. For example, the correlation between the B&L (1993) and NSD series over the whole

⁵ Behrman and Rosenzweig (1994) discuss some of the shortcomings of UNESCO's educational data.

sample stands at a respectable 0.81. An examination of average values for different geographic regions and of their evolution over time also reveals a fairly consistent and reasonable pattern. Industrialized countries and socialist economies display much higher attainment rates than less developed countries. Within this last group, Africa lies at the bottom of the distribution, while Latin America does fairly well and Southeast Asia features the largest improvement over the period.

This high overall correlation, however, hides significant discrepancies between the two data sets, both over time and across countries. Figure 1 shows B&L's (1996) and NSD's estimates of the average years of total schooling of the population over 15 for OECD countries in 1985. The correlation for the 23 countries in this sample (there are no data for Luxembourg) is now 0.574, but drops to zero (0.063) when we exclude the four countries with the lowest levels of schooling. When we disaggregate, the correlation is fairly high at the university level (0.767) and much lower for primary (0.362) and secondary (0.397) attainment.

Figure 1: Average years of schooling in 1985: B&L (1996) vs NSD





Notes:

- The estimates refer to the population over 15 in the case of Barro and Lee and to the age group 15-64 in Nehru et al.

- The estimated equation is of the form $H.NSD = 4.50 + 0.503 H.B\&L$, $t = 3.21$, $R^2 = 0.329$. The flattest line in the figure is the regression line fitted after excluding the four countries with the lowest schooling levels. The thinnest and steepest line is the "diagonal", where all the observations would fall if both sources agreed.

- Legend: Tu = Turkey; Por = Portugal; CH = Switzerland; Sp = Spain; Aus = Australia; It = Italy; Be = Belgium; Ge = West Germany; NL = Netherlands; Fr = France; NZ = New Zealand; Gr = Greece; Ost = Austria; Is = Iceland; Dk = Denmark; Nor = Norway; Fin = Finland; Swe = Sweden; Can = Canada; UK = United Kingdom; Jap = Japan; USA = United States; Ir = Ireland.

Figure 2: Average years of schooling by level in the OECD: B&L (1996) vs. NSD



- Notes:*
- Unweighted averages over the available OECD countries. Neither source reports data for Luxembourg. The sample excludes New Zealand except for average total years of schooling, as NSD only provide data on this variable but not its breakdown by level.
 - The data are for the age group 15-64 in the case of NSD and for the population aged 15 and over in Barro and Lee (1996).
 - The last year for the NSD series is 1987, rather than 1990.

Figure 2 shows the evolution of average years of total schooling (H) in the average OECD country and the breakdown of H by (primary, secondary and tertiary) levels ($H1$, $H2$ and $H3$) according to the same two sources (B&L, 1996 and NSD). If we focus on average years of total schooling, both data sets display an increasing trend, although it is much more marked in the case of B&L. In terms of their levels, NSD's figures on average attainment are significantly higher, although the difference between the two sets of estimates diminishes over time and becomes minor towards the end of the period. In principle, this discrepancy may be due at least in part to the difference between the age groups considered in the two studies. While B&L focus on the population aged 15 and over, NSD attempt to measure the educational attainment of the 15 to 64 age group. Since the older cohorts included in the B&L sample and excluded by NSD are typically less educated than the rest of the population, we would expect Barro and Lee's attainment estimates to be somewhat lower than Nehru et al's.

Significant differences between the two sources emerge when we disaggregate by educational level. In terms of secondary schooling the trend is quite similar in both cases but NSD's estimates are, unexpectedly, lower on average than B&L's. At the primary level, NSD's attainment figures are implausibly high, exceeding the duration of this school cycle (which is around six years on average), and display a downward trend. This "finding" that primary schooling levels have decreased over time in industrial countries is extremely suspicious, for it implies that new entrants into the labour force have less primary schooling than the older generations -- in spite of the rapid increase of enrollment rates observed over the relevant period.

For OECD countries we have some alternative sources that can be used to assess the likely accuracy of the B&L and NSD series. In particular, the OECD has published some reasonably complete educational statistics for most of its member countries. Although these data refer only to the last few years, and are therefore not an alternative to the other sources for the statistical analysis of the impact of education on growth, a comparison of the three sets of figures may perhaps give us some clues as to the possible shortcomings of the B&L (1996) and NSD data sets.

Table 2 summarizes the most relevant data. Notice that although both the year and the age groups differ somewhat across the three sources (see the notes to the table), the figures should be roughly comparable. The breakdown by educational level is also comparable with the one used by Barro and Lee (1996), although the OECD provides more detail. In particular, this last source disaggregates secondary attainment into two levels and, for most countries, reports data on advanced vocational programmes⁶ separately.

The differences across the various sources are quite significant. On the whole, the picture which emerges from the OECD figures seems to be the more plausible one -- at least in the sense of conforming better to common perceptions as to the relative educational levels of different countries. As for the other two sources, both contain rather implausible features and it is difficult to choose

⁶ These programmes generally correspond to level 5 of the International Standard Classification of Educational Attainment Levels (ISCED5).

between them. Starting with the relative positions of different countries in terms of average total years of schooling (reported in the last three columns of the table),⁷ we find a number of large discrepancies. Barro and Lee's estimates for Austria, France, Norway and Portugal are much lower than those given in the other sources, while their figure for New Zealand is much higher. On the other hand, NSD give very low figures for Australia, Switzerland and Germany, an extremely high estimate for Ireland (which is probably an error) and an implausibly high number for Greece.⁸ The overall correlation with the OECD estimates is higher for Barro and Lee (0.807) than for NSD (0.531) but this is due to a large extent to the Irish outlier.

In the case of Barro and Lee, it is possible to make a detailed comparison by levels of schooling with the OECD data that may give us some idea of the likely sources of some of their more implausible results.⁹ We observe that OECD estimates of secondary attainment are generally higher than Barro and Lee's.¹⁰ The difference exceeds forty points in Austria, Germany, Finland, Denmark, Norway and the UK, and is quite important in a number of other European countries and in Japan. We think the main reason for these large discrepancies has to do with the treatment of apprenticeships and other vocational training programmes, which are included in the OECD data but probably not in the UNESCO series used by Barro and Lee. Differences in tertiary attainment are significant as well and also seem to be related to the treatment of (higher-level) vocational programmes. In particular, Barro and Lee seem to report ISCED 5 studies (see footnote 6) as part of university schooling but, even accounting for this, significant differences remain in some cases.

Turning from the cross-section to the time-series dimension of the data, another disturbing feature of the human capital series is the existence of sharp breaks and implausible changes in attainment levels over very short periods. This problem affects the B&L data set much more than the NSD series, which are much smoother essentially by construction. Figures 3 and 4 show the evolution of Barro and Lee's (1996) secondary and university attainment rates for the population over 25 in a number of countries that display extremely suspicious patterns. In all cases, the sharp break in the series signals in all probability a change of criterion in the elaboration of educational statistics. Similar inconsistencies are present in other countries as well.


⁷ To estimate the average years of schooling on the basis of the OECD data we have used the following durations: Primary, 6 years; Secondary I, 9 years; Secondary II, 12 years; ISCED 5, 14 years; ISCED 6 and 7, 16 years. Since the computation assumes that everybody who started a certain level has completed it, the resulting figures should overstate the true years of schooling but, hopefully, not so much the relative positions of the different countries, which is what we are trying to get at. Our comparisons are based on the standardized attainment figures shown in Table 2, which are constructed by normalizing each estimate of average years of schooling by the unweighted average of the available contemporaneous observations in each data set.

⁸ According to NSD the average years of primary schooling in Ireland ranged between 15 in 1960 and just over 11 in 1985. Both figures are much higher than those for any other country and are of the order of twice the duration of this level of schooling. Greece does not appear in Table 2 because the OECD reports no data for this country. Greece is ranked by NSD ahead of Switzerland, Australia, Belgium, the Netherlands and France.

⁹ For a comparison of Barro and Lee's latest series with more recent OECD data, see Barro and Lee (2000).


¹⁰ Original OECD figures add up to 100% when we sum primary, secondary and tertiary attainment rates. Since this implies that everybody has received some schooling, we have corrected the figures using Barro and Lee's estimate of the fraction of the population with no schooling. The table reports the original primary attainment figure minus the no schooling fraction from Barro and Lee (1996).

Figure 3: Evolution of university attainment levels, Australia, New Zealand and Canada



- Source: Barro and Lee (1996). Population aged 25 and over.

Figure 4: Evolution of secondary attainment levels, Netherlands, New Zealand and Canada



- Source: Barro and Lee (1996). Population aged 25 and over.

The preceding discussion is far from providing an exhaustive list of the suspect features of different educational data sets. On the other hand, it is probably sufficient to conclude that --despite the fact that the contributions we review represented a significant advance in this area-- the data on human capital stocks that have been used in most empirical analyses of the determinants of growth are still of dubious quality. Remaining problems are probably due in part to the fact that the primary statistics used in these studies are not consistent, across countries or over time, in their treatment of

vocational and technical training and other courses of study,¹¹ and reflect at times the number of people who have started a certain level of education and, at others, those who have completed it. Additional problems may be traced to the procedure used in the construction of the data and even to computational mistakes. Thus, NSD's neglect of census data probably accounts for their unreasonable results in terms of the overall level and trend of primary and secondary schooling while Barro and Lee's (1993) approximation to a perpetual inventory method is probably not entirely satisfactory.

Concerns about poor data quality and about its implications for empirical estimates of the growth effects of human capital have motivated some recent studies that attempt to increase the signal to noise ratio in the schooling series by exploiting additional sources of information and introducing various corrections. As we have seen, the latest version of the Barro and Lee data set (B&L, 2000) incorporates various refinements of the procedure used to fill in missing observations. De la Fuente and Doménech (D&D, 2000) and Cohen and Soto (C&S, 2000), on the other hand, have focused on trying to clean up the available census and survey data.¹² In the following section we will discuss the construction of attainment series that update the work in our previous paper, and we will carry out a comparative analysis of these three sources in a sample of OECD countries.

3. Educational attainment in the OECD: A revised set of estimates for 1960-95

As we have seen in the previous section, the available schooling data contain a large amount of noise that can be largely traced back to inconsistencies of the underlying primary statistics. In an attempt to reduce this noise, we have constructed educational attainment series for the adult population of a sample of 21 OECD countries covering the period 1960-95.¹³ The complete series and a detailed set of country notes are included in Appendix 1.

This data set is a revised and partially extended version of the one described in de la Fuente and Doménech (D&D, 2000). There are two major changes relative to our earlier estimates. First, we have incorporated a fair amount of new information supplied by the national statistical offices of around a dozen countries in response to a request by the OECD's Statistics and Indicators Division that was accompanied by our previous paper. Second, we have extended the series to 1995 for around three fourths of the sample using national sources, the latest editions of *Education at a Glance* (EAG) and Labour Force Survey data supplied directly by the OECD. For the remaining countries, these sources

¹¹ Steedman (1996) documents the existence of important inconsistencies in the way educational data are collected in different countries and argues that this problem can significantly distort the measurement of educational levels. She notes, for example, that countries differ in the extent to which they report qualifications not issued directly (or at least recognized) by the state and that practices differ as to the classification of courses which may be considered borderline between different ISCED levels. The stringency of the requirements for the granting of various completion degrees also seems to vary significantly across countries.

¹² Cohen and Soto (2001) construct a schooling data set for a sample of 95 countries covering the period 1960-2000 at 10-year intervals. They collect census and survey data from UNESCO, the OECD's in-house educational data base and the websites of national statistical agencies, and exploit to the extent possible the available information on attainment levels by age group to fill in missing cells through forward and backward extrapolations. Remaining gaps in the data are filled using enrollment rates from UNESCO and other sources. Their estimates refer to the 15 to 64 age group.

¹³ Iceland, Luxembourg, Turkey and recent OECD members are left out because of the scarcity of information

do not seem to be consistent with the national data on which our series are primarily based (at least in part because of differences in the age groups covered by them), so an extension will have to wait until the results of the current round of censuses are available. Finally, there are a few minor changes in the series that are mostly due to corrections of i) arithmetical mistakes, ii) the dates attributed to some census observations (usually by one year) and iii) changes in the assumed duration of some schooling cycles (following the indications of national statistical offices or to make them consistent with the cutoff levels we have used for some countries).

Table 3: Attainment levels and codes

<i>code</i>	<i>level</i>
<i>L0</i>	Illiterates
<i>L1</i>	Primary schooling
<i>L2.1</i>	Lower secondary schooling
<i>L2.2</i>	Upper secondary schooling
<i>L2</i>	Total secondary schooling = $L2.1 + L2.2$
<i>L3.1</i>	Higher education, first cycle or shorter courses
<i>L3.2</i>	Higher education, second cycle or full-length courses
<i>L3</i>	Total higher education = $L3.1 + L3.2$

We aim to provide estimates of the fraction of the population aged 25 and over that has started (but not necessarily completed) each of the levels of education shown in Table 3 (illiterates (*L0*), primary schooling (*L1*), lower and upper secondary schooling (*L2.1* and *L2.2*) and two levels of higher education (*L3.1* and *L3.2*)). For some countries, however, the available data may refer to a different age group or to those who have completed each schooling level, and it is not always possible to detect when this is the case. We have tried (with uncertain success) to include higher vocational training (ISCED 5) in the first level of higher attainment. We report *L0* only for the four countries where illiteracy rates are significant during the sample period (Portugal, Greece, Spain and Italy). For the rest of the sample, the lowest reported category is *L1*, and it includes all those who have not reached secondary school.

Our approach has been to collect all the information we could find on educational attainment in each country, both from international publications and from national sources (census and survey results, statistical yearbooks and unpublished national and OECD data), and use it to try to reconstruct a plausible attainment profile for each country. For those countries for which reasonably complete series are available, we have relied primarily on national sources. For many of the rest, we start from the most plausible set of attainment estimates available around 1990 or 1995 (taken generally from OECD sources) and proceed backwards using all the assembled information and trying to avoid unreasonable jumps in the series by choosing the most plausible figure when several are available for the same year, and by reinterpreting some of the data (as referring to broader or narrower schooling categories than the reported one) when it seems sensible to do so. Missing

observations are then filled in a variety of ways. Where possible, we interpolate between available observations. Otherwise, we use information on educational attainment by age group in order to make backward or forward projections, or rely on miscellaneous information from a variety of sources in order to construct estimates of attainment levels. We have avoided the use of flow estimates based on enrollment data because they seem to produce implausible time profiles.

Clearly, the construction of our series involves a fair amount of guesswork. Our "methodology" looks decidedly less scientific than the apparently more systematic estimation procedures used by other authors starting from supposedly homogeneous data. As discussed in the previous section, however, even a cursory examination of the data shows that there is no such homogeneity. Hence, we have found it preferable to rely on judgment to try to piece together the available information in a coherent manner than to take for granted the accuracy of the primary data. As we will see below, the results do look more plausible than most existing series, at least in terms of their time profile, and perform rather well in terms of a statistical indicator of data quality.

A number of countries do not separate primary education from lower secondary schooling and report a single attainment level that comprises all mandatory courses. To preserve the homogeneity of our attainment categories, we have estimated the breakdown of compulsory schooling into *L1* and *L2.1*. For some countries we have managed to find enough information to make what seemed to be a reasonable guess. For others, we have used data from close neighbours. In particular, we have used information for the US, Germany and Norway to estimate the breakdown in Canada, Austria and Denmark, respectively.¹⁴ Finally, for those countries for which there is no obvious candidate for this role (the UK and Japan), we have used an ad-hoc regression estimate of the relevant ratio. Using those countries in the sample for which there is decent direct information on *L1* and *L2.1*,¹⁵ we estimate the following equation with pooled data:

$$(1) \quad L2.1/(L1+L2.1) = 0.1730 + 0.00947(L3+L2.2) + 0.08272(L3/L2.2) - 0.025*trend \quad \text{adj. } R^2 = 0.674$$

$$\quad \quad \quad (1.71) \quad \quad (16.23) \quad \quad (2.03) \quad \quad (1.77)$$

where the numbers in parentheses below each coefficient are *t* ratios. That is, we hypothesize that those countries that are more "efficient" in getting students into the upper schooling cycles will also have greater accession rates to lower secondary schooling. Hence we specify the weight of lower secondary attainment relative to primary attainment as a function of university and upper secondary attainment and the ratio of the two, and allow it to vary systematically over time. Since the fit of the equation is reasonably good, we use it to estimate the lower-secondary/compulsory attainment ratio in those countries for which this information is not available.

Using our attainment series, we construct an estimate of average years of total schooling for each country and period. The assumed cumulative duration of the different school cycles in each country is shown in Table 4. In constructing these series we are implicitly assuming that everybody who starts a

¹⁴ See the country notes in Appendix 1 for further details.

¹⁵ We exclude from the sample used in this regression Norway, Finland, New Zealand and Germany (as well as Denmark, Austria and Canada).

given school cycle does eventually complete it, which is clearly not the case. Hence, our figures will be biased upward and are not strictly comparable with Barro and Lee's (1996 or 2000) average schooling series, which do make use of estimates of completion rates.^{16,17}

Table 4: Cumulative years of schooling by educational level

	<i>L1</i>	<i>L2.1</i>	<i>L2.2</i>	<i>L3.1</i>	<i>L3.2</i>
Australia	7	11	13	15	16
Austria	4	8/9 ^a	13	15	17
Belgium	6	9	12	15	16
Canada	6	9	12	15	16
Denmark	6	9	13	14	17
Finland	6	9	12	14	17
France	5	9	12	14	16
Germany	4	10	13	15	17
Greece	6	9	12	16	16
Ireland	6	9	12	14	16
Italy	5	8	13	15	18
Japan	6	9	12	14	16
Netherlands	6	10	12	16	17
Norway	6	9	12	14	16
New Zealand	6	11	13	15	16
Portugal	6	8	12	14	16
Spain	5	8	12	14	17
Sweden	6	9	12	14	16
Switzerland	7 ^b	9	13	16	17
UK	6	9	12	14	16
USA	4 ^c	8 ^c	12	14	16
Mode	6	9	12	14	16

- Sources: Education at a Glance 1997 (OECD, 1998), except figures in bold type (WDI, World Bank, 1999) and in italics (national sources).

^(a) The duration of compulsory schooling in Austria changed from 8 to 9 years in the mid sixties.¹⁸

^(b) The duration of *L1* in Switzerland is 6 years, but the cutoff between *L1* and *L2.1* in our data is 8. Hence, we use the average of these two figures to calculate the average years of schooling.

^(c) Set in accordance with our cutoffs for *L1* and *L2.1*.

Data availability varies widely across countries. Table 5 shows the fraction of the reported data points that correspond to "direct observations" and the earliest and latest such observations available for secondary and higher attainment levels. The number of possible observations is typically either 21 or 24 for each level of schooling depending on whether the series ends in 1990 or 1995 (two

¹⁶ The available data on completion rates display the same anomalies we have discussed above in connection with attainment and enrollment rates.

¹⁷ Except for the fact that they refer to different age groups, our estimates of years of schooling do seem to be comparable to Cohen and Soto's (2001). These authors, however, do not provide sufficient details to identify any possible differences in the procedures followed in the two studies.

¹⁸ We allow for this in a crude way when computing the average years of schooling in Austria. For this calculation, we assume that the cumulative duration of *L2.1* is 8 years until 1975 (since those completing *L2.1* in 1965 will not enter the 25+ age group until that year). For later years, we attribute to *L2.1* a cumulative duration which is a weighted average of 8 and 9 years, with the weight of the second figure set equal to the fraction of the population 25 and over in the current year that studied under the new system. The data on the age structure of the population required for the adjustment was obtained from the UN's *Demographic Yearbook* (DYB).

sublevels and a total¹⁹ times seven or eight quinquennial observations). In the case of Italy, there seem to be no short higher education courses, so the number of possible observations at the university level drops to eight. We count as direct observations backward and forward projections constructed using detailed census data on educational attainment by age group and the age structure of the population, and various "reasonable guesses" that incorporate some information from census or survey data.

As can be seen in the Table, for most of the countries in the sample we have enough primary information to reconstruct reasonable attainment series covering the whole sample period. The more problematic cases are highlighted using bold characters. In the case of Italy, the main problem is that much of the available information refers to the population over six years of age. For Denmark and Germany (at the secondary level), the earliest available direct observation refers to 1970 or later. In these two cases, we have projected attainment rates backward to 1960 using the attainment growth rates reported in OECD (1974), but we are unsure of the reliability of this extrapolation.

Table 5: Availability of primary data

	<u>secondary attainment</u>			<u>university attainment</u>		
	direct/tot. observ.	first observ.	last observ.	direct/tot. obs.	first observ.	last observ.
USA	24/24	1960	1995	24/24	1960	1995
Netherlands	12/24	1960	1995	12/24	1960	1995
Italy	15/24	1961	1999	5/8	1960	1998
Belgium	13/24	1961	1995	12/24	1960	1995
Spain	12/21	1960	1991	12/21	1960	1991
Greece	15/24	1961	1995	15/24	1961	1997
Portugal	12/21	1960	1991	8/21	1960	1991
France	12/21	1960	1989	12/21	1960	1990
Ireland	15/24	1961	1998	11/24	1961	1998
Sweden	9/24	1960	1995	9/24	1960	1995
Norway*	15/24	1960	1998	9/24	1960	1998
Denmark*	9/24	1973	1994	12/24	1973	1994
Finland*	16/24	1960	1995	21/24	1970	1995
Japan*	8/21	1960	1990	12/21	1960	1990
N. Zealand	10/24	1965	1998	10/24	1965	1998
UK*	6/21	1960	1993	10/21	1960	1991
Switzerland	15/24	1960	1995	15/24	1960	1995
Austria*	11/24	1961	1995	7/24	1961	1995
Australia	11/24	1965	1997	11/24	1966	1997
W. Germany	11/24	1970	1995	17/24	1961	1995
United Germany	6/6	1991	1995	6/6	1991	1995
Canada*	15/24	1961	1996	21/24	1960	1996

- (*) Countries where primary and lower secondary attainment are generally not reported separately.


¹⁹ For those countries where primary and lower secondary attainment are generally not reported separately (identified by an asterisk in Table 5), the three categories reported in the column for secondary attainment are L2.1, L2.2 and L1+L2.1 (rather than L2.1, L2.2 and L2 as in the other cases).

3.1 An example: the case of higher education in Canada

To give the reader a flavour for the way our series have been constructed and for their limitations, we will discuss in some detail the case of higher education in Canada. This is a country for which there is a considerable amount of published information that displays, if taken literally, a rather implausible pattern. It is also a case in which the new data supplied by national statistical agencies allows us to check our original estimates (D&D 2000).

The essence of our approach is captured by Figure 5. The thickest line in the figure (labeled *B&L.96*) describes Barro and Lee's (1996) higher educational attainment series for the population aged 25 and over, which is based on Unesco and UN data. The implausible hump-shaped pattern of the series strongly suggests that the 1975 and 1980 observations refer to a broader concept of higher attainment than the rest of the data. Our (2000) estimate was based on the guess that, unlike the rest of the observations, the data for these two atypical years included upper-level vocational training courses. Using other available information, we tried to homogenize the series by consistently including or excluding an estimate of this category, obtaining the more plausible profile described by the two thinnest lines shown in the figure.²⁰ The higher of these lines (*L3.00*) refers to higher education in a broad sense, and the lower one (*UNIV*) to strictly university attainment (which is not exactly the same thing as *L3.2*). The dots lying on these two lines represent actual data taken from various sources and attributed to the exact year to which they correspond (and not to the closest multiple of five). For the rest of the years, the series are completed through linear interpolation.

Figure 5: University attainment in Canada, B&L (1996) vs. D&D (2000) and this paper



²⁰ The details can be found in de la Fuente and Doménech (2000).

Our current estimates correspond to the lines of intermediate thickness shown in the figure and are based mainly on data supplied directly by Statistics Canada. The series labeled *L3.2.02* (completed Bachelor's or higher) extends back to 1951 and displays a plausible profile. To recover total tertiary attainment (*L3.02*) we add to the previous series an estimate of *L3.1*. From 1971 onward, this estimate is also taken directly from Statistics Canada. For the previous period, we use other sources to estimate this category as discussed in the country notes (see section 1 of Appendix 1). The new data suggest that our 2000 estimate was essentially correct for the second part of the sample period, but probably overestimated higher attainment until 1975.

3.2. A comparison of three recent data sets

In this section we compare our attainment series with those constructed by Cohen and Soto (2001) and with the latest version of the Barro and Lee data set (B&L 2000). We find that there are significant differences across these three sources in terms of both their cross-section and their time profiles. In the cross-section dimension, our estimates tend to lie between the other two. In the time dimension, both our series and Cohen and Soto's display considerably less volatility than Barro and Lee's. A detailed country by country comparison of these three attainment series and a set of pairwise scatterplots for the data in levels and in growth rates can be found in Figures A.1-A.6 in Appendix 1.

**Table 6: Normalized average years of schooling over 1960-90
Correlation across data sets**

	B&L	C&S	D&D
B&L (2000)	1.000		
C&S (2001)	0.885	1.000	
D&D (2002)	0.925	0.924	1.000

- Note: Data from the last 3 columns of Table 7.


To compare their cross-section profiles, we begin by normalizing each of the years of schooling series by its contemporaneous sample mean.²¹ Table 7 shows the resulting attainment indices for 1960 and 1990. The last three columns of the table refer to average attainment calculated over the entire period 1960-90. Focusing on this last set of figures, we observe that although the correlation across data sets is quite high (see Table 6), there are also large discrepancies among them. This is illustrated in Figure 6, which shows the differences in normalized average attainment across sources, taking as a reference Barro and Lee's estimates. To construct a rough measure of the degree of agreement across series, we will say that two sources (or the three of them) agree for a given country if the (maximum) difference between them in terms of normalized years of schooling is less than 5% of their average value. We find that there is not a single country for which all three sources agree. The number of

²¹ As has already been noted, our average years of schooling series is not directly comparable with Barro and Lee's.

Table 7: Normalized average years of schooling

	1960			1990			average 1960-90		
	B&L00	D&D02	C&S01	B&L00	D&D02	C&S01	B&L00	D&D02	C&S01
USA	129.3	126.3	126.1	135.2	119.1	115.5	136.5	123.4	120.6
Australia	140.8	117.7	121.7	114.1	121.1	116.8	128.2	122.4	119.5
New Zealand	142.7	125.1	111.3	126.0	113.8	100.8	136.6	119.2	105.8
Switzerland	109.0	124.8	135.8	111.8	114.9	118.6	113.7	120.0	125.7
Canada	125.0	124.1	112.9	118.3	119.7	113.1	122.7	122.4	113.2
West Germany	123.6	118.5	117.9	102.1	121.7	120.9	109.8	121.1	121.3
Denmark	133.6	129.0	112.5	114.2	110.2	105.6	119.6	119.3	108.8
Norway	91.2	115.8	112.1	122.3	104.4	112.7	102.2	109.3	112.7
Japan	102.6	103.1	117.4	103.9	105.6	109.2	101.1	104.6	112.1
Sweden	114.2	96.2	107.5	107.9	99.8	110.2	110.7	97.0	109.4
UK	114.5	102.5	112.9	98.5	98.9	112.4	104.0	100.1	112.9
Austria	100.2	107.7	102.6	92.6	106.3	100.1	97.5	105.1	101.2
Netherlands	78.7	97.0	103.3	97.0	102.9	98.1	95.1	99.6	100.9
Belgium	111.4	92.5	91.6	95.0	94.7	91.8	104.0	94.1	91.1
Finland	80.2	91.5	84.9	106.8	103.1	98.2	94.3	98.5	91.3
France	86.3	97.3	83.4	85.2	98.2	94.8	84.2	99.2	89.8
Ireland	96.3	88.0	89.8	95.8	88.4	87.2	93.3	87.1	88.0
Greece	69.3	66.5	73.6	86.3	74.3	79.7	77.5	70.5	75.9
Italy	68.1	64.7	72.1	69.4	75.6	83.3	69.0	70.1	77.3
Spain	54.3	59.5	71.7	68.6	66.7	77.2	61.7	60.6	73.5
Portugal	29.0	52.3	39.0	48.8	60.2	54.1	38.3	56.3	48.9
average	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
average in years	6.70	8.36	8.07	8.87	10.64	10.93	7.67	9.47	9.59

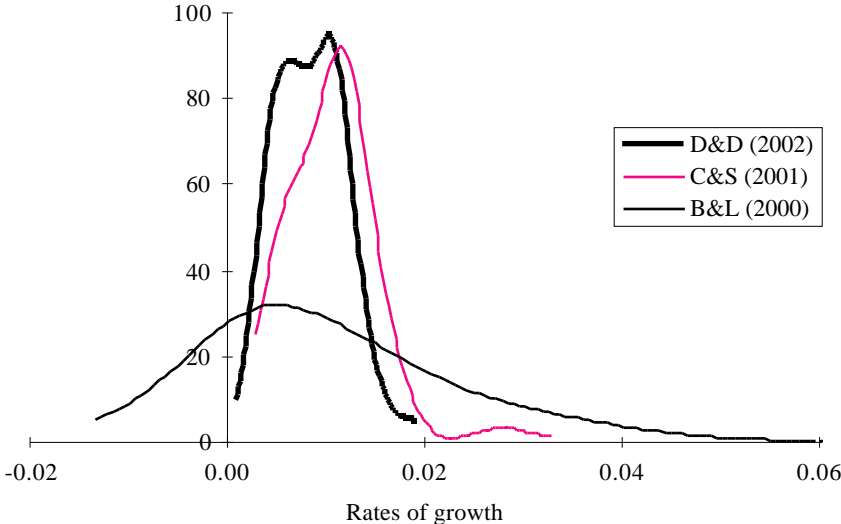
Figure 6: Normalized years of schooling, differences with B&L (2000) based on average schooling over the entire sample period



observations for which there is agreement is 9 (out of 21) when our data are compared with either the B&L or C&S series, and only 4 when these two series are compared with each other.


When we turn to the time profiles, both our series and Cohen and Soto's display a considerably smoother and more plausible pattern than the Barro and Lee data set. Barro and Lee's series contain a large number of sharp breaks that give a distorted image of the pattern of human capital accumulation and may obscure its relationship with productivity growth. This is clearly illustrated in Figure 8, where we have plotted the fitted distribution of the annualized growth rate of average years of schooling for all countries and years in each data set. The difference in the range of this variable across data sets is enormous: while our annual growth rates range between 0.09% and 1.92%, and those of Cohen and Soto between 0.27% and 3.27%, Barro and Lee's go from -1.35% to 6.13%; moreover, 19% of the observations in this last data set are negative, and 16.7% of them exceed 2%.

Figure 7: Fitted distribution of the growth rate of years of schooling, Barro and Lee (2000), Cohen and Soto (2001) and this paper



As shown in Figure 8, an implausibly broad range of values (for the data in growth rates) is a common feature of all versions of the Barro and Lee data set. We believe that this anomaly, which seems to arise from these authors' reliance on UNESCO data, cannot be corrected by any improvements in the fill-in procedure alone. We also suspect that this feature may explain why these series often generate implausible results in growth regressions, particularly when they are estimated using panel or differenced specifications. A first indication of this is that the coefficient of a univariate regression of the growth rate of output per employed worker on the growth rate of years of schooling (with both variables measured in deviations from their contemporaneous sample averages) increases from 0.173 (with a t -ratio of 1.89) with the Barro and Lee (2000) data to 0.348 (with $t = 3.07$) with Cohen and Soto's series, and to 1.03 (with $t = 3.16$) with those constructed in this paper. The results we report later in the paper are also consistent with this hypothesis.

Figure 8: Fitted distribution of the growth rate of years of schooling, different versions of the Barro and Lee data set



One conclusion we draw from the above comparisons is that a fair amount of detailed work remains to be done before we can say with some confidence that we have a reliable and detailed picture of educational achievement levels in industrial countries or of their evolution over time. On the other hand, the existing data bases may contain sufficient information to provide reasonable statistical estimates of the contribution of human capital accumulation to productivity growth, particularly if we exploit differences in quality across them to correct for measurement error. This will be the objective of section 5 of the paper.

3.3 Measuring data quality: preliminary estimates of reliability ratios

In this section we will construct an indicator of the quality of a number of schooling series using an extension of the procedure suggested by Krueger and Lindhal (2001). As these authors note, the information content of a noisy proxy can be measured by its *reliability ratio*, defined as the ratio of signal to signal plus measurement noise in the data. When several noisy measures of the same magnitude are available, estimates of their respective reliability ratios can be obtained by regressing these variables on each other. Under certain assumptions, the coefficients obtained in this manner can be used to approximate the bias induced by measurement error (which will be a decreasing function of the reliability ratio) and to obtain consistent estimates of the parameters of interest in growth regressions.

Let H be the true stock of human capital and let $P_1 = H + \varepsilon_1$ be a noisy proxy for this variable, where the measurement error term ε_1 is an iid disturbance with zero mean and uncorrelated with H . The reliability ratio of this series (r_1) is defined as

$$(2) r_1 \equiv \frac{\text{var } H}{\text{var } P_1} = \frac{\text{var } H}{\text{var } H + \text{var } \varepsilon_1}$$

Assume now that in addition to P_1 we have a second imperfect measure of human capital, $P_2 = H + \varepsilon_2$, where ε_2 is also iid noise. Then, the covariance between P_1 and P_2 can be used to obtain an estimate of the variance of H whenever the measurement error terms ε_1 and ε_2 are uncorrelated. Under this assumption, r_1 can be estimated by

$$(3) \hat{r}_1 = \frac{\text{cov}(P_1, P_2)}{\text{var } P_1}$$

which happens to be the formula for the OLS estimator of the slope coefficient of a regression of P_2 on P_1 . Hence, to estimate the reliability of P_1 we run a regression of the form $P_2 = c + r_1 P_1$.²² Notice, however, that if the measurement errors of the two series are positively correlated ($E\varepsilon_1\varepsilon_2 > 0$) as may be expected in many cases, \hat{r}_1 will overestimate the reliability ratio and hence understate the extent of the attenuation bias induced by measurement error.

We will build on this approach to construct indicators of the information content of the most widely used cross-country schooling data sets, restricting ourselves to the sample of 21 OECD countries covered by our series. For now, we will maintain Krueger and Lindhal's assumption that measurement error is uncorrelated across data sets (and with other variables of interest), and we will exploit the availability of a number of alternative human capital series to construct a minimum-variance estimator of the reliability ratio. In section 5 we will relax this assumption and allow for correlated measurement errors.

Proceeding as indicated above, for each data set P_j we first construct a series of "pairwise" reliability ratio estimates, \hat{r}_{jk} , by using P_j to try to explain alternative schooling series P_k ; that is, for each j we estimate a system of equations of the form

$$(4) P_k = c_k + r_{jk} P_j + u_k \quad \text{for } k = 1, \dots, K$$

where k denotes the "reference" data set and varies over the last available version of all schooling series different from j . The reliability ratio of Barro and Lee's (2000) data set, for instance, is estimated by using these authors' estimate of average years of schooling as the explanatory variable in a set of regressions where the reference (dependent) variables are the average years of schooling estimated by Kyriacou (1991), NSD (1995), Cohen and Soto (2001) and ourselves. Other versions of the Barro and Lee data set, however, are not used as a reference because the correlation of measurement errors across the same family of schooling series is almost certainly very high and this will artificially inflate the estimated reliability ratio.

Under the assumption that measurement error is uncorrelated across families of data sets (i.e. that $E\varepsilon_j\varepsilon_k = 0$ for $j \neq k$ when j and k belong to different families) all the pairwise estimates of r_j obtained above will be consistent and so will be any weighted average of them,

$$(5) \bar{r}_j = \sum_k \omega_k \hat{r}_{jk} \quad \text{where } \sum_k \omega_k = 1.$$

²² Intuitively, regressing P_2 on P_1 gives us an idea of how well P_1 explains the true variable H because measurement error in the dependent variable (P_2 in this case) will be absorbed by the disturbance without generating any biases. Hence, it is almost as if we were regressing the true variable on P_1 .

To obtain the most efficient estimator of r_j , we choose the weights ω_k in (5) so as to minimize the variance of \bar{r}_j . The resulting estimator, which will be denoted by \hat{r}_j , can be implemented by imposing a common slope coefficient across the equations in (4) and estimating the system as a restricted SUR.²³ Hence, we will refer to \hat{r}_j as the *SUR reliability ratio*.

The exercise we have just described is repeated for several transformations of average years of schooling.²⁴ In particular, we estimate reliability ratios for years of schooling measured in levels (H_{it}) and in logs (h_{it}), for average annual changes in both levels and logs measured across successive (quinquennial) observations (ΔH_{it} and Δh_{it}), for log years of schooling and for their annualized changes measured in deviations from their respective country means ($h_{it} - h_i$ and $\Delta h_{it} - \Delta h_i$) and for average annual log changes computed over the period 1965-85 (Δh_i).²⁵ Notice that Δh_{it} corresponds to annual growth rates and that $h_{it} - h_i$ and $\Delta h_{it} - \Delta h_i$ are the "within" transformations often used to remove fixed effects. We also estimate all the reliability ratios (except for the case of Δh_i) twice, once with the raw data and a second time after removing the period means from the different schooling series. Since all our growth estimates (reported below) include fixed period effects, this second set of reliability ratios is the relevant one for the analysis of the sensitivity of human capital estimates to data quality we will carry out in a later section.

The results are shown in Tables 8a and 8b with the different data sets arranged by decreasing average reliability ratios. The last row of each table shows the average value of the reliability ratio for each type of data transformation (taken across data sets), and the last column displays the average reliability ratio of each data set (taken across data transformations). Our mean estimate of the reliability ratio in the OECD sample is 0.386 for the raw data and 0.335 after removing period fixed effects. This suggests that the average estimate of the coefficient of schooling in a growth equation is likely to suffer from a substantial downward bias, even without taking into account the further loss of signal that arises when additional regressors are included in these equations (see below). The bias will tend to be smaller when the data are used in levels or logs, even when fixed effects are removed, but is likely to be extremely large in specifications that use data differenced over relatively short periods. The average reliability ratio is only 0.227 for the data in quinquennial log differences, and 0.075 for level differences taken at the same frequency.²⁶ It should be noted that, while reliability ratios must lie between zero and one, some of the estimates reported in Table 8 fall outside these bounds. We interpret these results as

²³ Actually, the restricted SUR yields an approximation to the desired estimator because the number of available data points differs across data sets and the SUR is estimated only over the common observations.

²⁴ Notice that each set of estimates will be based on assumptions about the properties of the measurement error term that are not necessarily consistent with each other.

²⁵ This is the longest period over which all the available schooling series overlap. For some countries, the earliest observation provided by Kyriacou (1991) corresponds to 1970 or 1975. Rather than dropping these observations, in these cases we set Δh_i equal to the average growth rate over the available period.

²⁶ Estimated reliability ratios are generally somewhat higher in larger samples (see de la Fuente 2002a). This is likely to be misleading, however, because the number of available primary sources that can be drawn upon to construct estimates of educational attainment is probably higher in developed than in underdeveloped countries. As a result, the variation across data sets is likely to be smaller in LDCs, and this will tend to raise the estimated reliability ratio. To a large extent, however, this will simply reflect a higher correlation of errors across data sets (i.e. an upward bias in the estimated reliability ratio).

Table 8: SUR estimates of reliability ratios, OECD 21 sample**a. Raw data**

	H_{it}	h_{it}	ΔH_{it}	Δh_{it}	$h_{it}-h_i$	Δh_i	$\Delta h_{it}-\Delta h_i$	average
<i>D&D (2002)</i>	0.806 [0.027]	0.859 [0.032]	0.293 [0.110]	0.716 [0.133]	1.166 [0.035]	1.028 [0.205]	0.163 [0.121]	0.719
<i>D&D (2000)</i>	0.775 [0.029]	0.848 [0.034]	0.057 [0.098]	0.477 [0.131]	1.139 [0.040]	0.813 [0.255]	0.040 [0.103]	0.592
<i>C&S (2001)</i>	0.793 [0.033]	0.889 [0.032]	0.227 [0.084]	0.389 [0.060]	0.735 [0.021]	0.545 [0.101]	0.124 [0.054]	0.529
<i>B&L (2000)</i>	0.768 [0.035]	0.644 [0.023]	0.011 [0.024]	0.046 [0.028]	0.488 [0.036]	0.255 [0.092]	0.002 [0.014]	0.316
<i>Kyr. (1991)</i>	0.698 [0.097]	0.601 [0.079]	0.027 [0.019]	0.064 [0.020]	0.326 [0.027]	0.164 [0.050]	0.025 [0.011]	0.272
<i>B&L (1996)</i>	0.628 [0.041]	0.562 [0.029]	0.014 [0.023]	0.042 [0.027]	0.466 [0.034]	0.125 [0.088]	0.013 [0.014]	0.264
<i>B&L (1993)</i>	0.587 [0.037]	0.476 [0.025]	0.010 [0.021]	0.033 [0.016]	0.286 [0.027]	0.109 [0.082]	-0.010 [0.008]	0.213
<i>NSD (1995)</i>	0.308 [0.061]	0.383 [0.060]	-0.037 [0.040]	0.045 [0.047]	0.510 [0.074]	0.179 [0.075]	-0.137 [0.048]	0.179
average	0.670	0.658	0.075	0.227	0.639	0.402	0.027	0.386

b. Data in deviations from period means

	H_{it}	h_{it}	ΔH_{it}	Δh_{it}	$h_{it}-h_i$	$\Delta h_{it}-\Delta h_i$	average
<i>D&D (2002)</i>	0.754 [0.026]	0.775 [0.031]	0.337 [0.091]	0.769 [0.118]	0.917 [0.092]	0.246 [0.098]	0.633
<i>C&S (2001)</i>	0.806 [0.036]	0.912 [0.034]	0.330 [0.096]	0.467 [0.063]	0.547 [0.049]	0.185 [0.073]	0.541
<i>D&D (2000)</i>	0.720 [0.027]	0.761 [0.032]	0.100 [0.084]	0.550 [0.119]	0.818 [0.102]	0.074 [0.085]	0.504
<i>Kyr. (1991)</i>	0.723 [0.111]	0.600 [0.092]	0.024 [0.017]	0.065 [0.020]	0.111 [0.024]	0.026 [0.011]	0.258
<i>B&L (2000)</i>	0.707 [0.034]	0.603 [0.022]	-0.018 [0.025]	0.045 [0.031]	0.178 [0.040]	-0.016 [0.015]	0.250
<i>B&L (1996)</i>	0.559 [0.040]	0.516 [0.028]	-0.017 [0.024]	0.039 [0.030]	0.146 [0.041]	-0.007 [0.014]	0.206
<i>B&L (1993)</i>	0.526 [0.036]	0.436 [0.024]	-0.019 [0.022]	0.029 [0.017]	0.121 [0.021]	-0.017 [0.008]	0.179
<i>NSD (1995)</i>	0.278 [0.051]	0.330 [0.050]	-0.021 [0.036]	0.066 [0.046]	0.095 [0.046]	-0.115 [0.042]	0.106
average	0.634	0.617	0.090	0.254	0.367	0.047	0.335

Notes:

- Standard errors in brackets below each estimate.

- Data are reported at 5-year intervals except by Cohen and Soto who do it at 10-year intervals. We use linear interpolation (with the data in levels) to complete these series prior to all calculations.

- The version we use of Barro and Lee (1993) is actually taken from Barro and Lee (1994). We do not know if the two data sets are identical or if there are minor differences between them.

- Panel a corresponds to the variables as originally measured. The estimates shown in panel b are obtained after removing the corresponding period means. This is done by introducing period dummies in equation (4).

an indication that measurement error is likely to be correlated across data sets, a possibility that will be investigated below.

Our results indicate that the importance of measurement error varies significantly across data sets, although their precise ranking depends on the data transformation that is chosen. Two of the datasets most widely used in cross-country empirical work, those by Kyriacou (1991) and Barro and Lee (various years), perform relatively well when the data are used in levels but, as Krueger and Lindhal (2001) note, contain very little signal when the data is differenced. Recent efforts to increase the signal content of the schooling data seem to have been at least partially successful. Taking as a reference the average reliability ratio for the (1996) version of the Barro and Lee data set (0.264 for the raw data and 0.206 after removing period means), the latest revision of these series by the same authors has increased their information content by 20% (21% after removing period means), while the estimates reported in Cohen and Soto (2001) and in this paper raise the estimated reliability ratio by 100% (162%) and 172% (207%) respectively. Our updated data set has the highest information content of all those available, particularly when the data are used in growth rates or after removing fixed country effects.

4. Comparative growth results for different data sets

Doubts about the accuracy of existing data sets must raise concerns about the validity of the findings of empirical studies based on them. Concerns about data quality, however, also admit an optimistic interpretation of these results. Since there are no reasons to suspect that the available data contain systematic biases that may lead us to overestimate the contribution of human capital to productivity, the fact that the empirical results are quite favourable in some cases in spite of the dubious quality of the data suggest that improvements in this regard should lead to clearer and more conclusive results about education's contribution to economic growth.

In this final part of the paper we will provide some evidence that this is indeed the case. In this section, we examine the performance of different schooling data sets in a number of standard growth specifications and find a clear pattern of human capital coefficients rising with increases in data quality as measured by estimated reliability ratios. In the following section, we will use this correlation to produce a "meta-estimate" of the relevant parameter that corrects for measurement error bias. As we will see, our results support the hypothesis that the lack of correlation between productivity growth and human capital accumulation reported in some recent studies may be due to data deficiencies and suggest that the return to investment in education is substantial.

We will assume that the educational attainment of employed workers (HE) is one of the inputs in a constant-returns Cobb-Douglas aggregate production function which we will write in intensive form

$$(6) q_{it} = a_{it} + \alpha z_{it} + b^* h e_{it}$$

where q_{it} is the log of output per employed worker in country i at time t , z the log of the stock of physical capital per employed worker, he the log of the average number of years of schooling of

employed workers and a_{it} the log of total factor productivity (TFP). One difficulty we face when trying to estimate (6) is that our human capital data (H) generally refer to the adult population rather than to employed workers. To get around this problem, we will hypothesize that HE increases with population attainment and decreases with the ratio of employment to the adult population (E), i.e. that

$$(7) \quad he_{it} = c^*h_{it} - d^*e_{it}$$

where all variables are measured in logarithms. Substituting (7) into (6) we obtain the reduced-form production function

$$(8) \quad q_{it} = a_{it} + \alpha z_{it} + \beta h_{it} - \varphi e_{it}$$

where

$$(9) \quad \beta = bc \quad \text{and} \quad \varphi = bd.$$

Hence, our use of population data is likely to introduce a bias in the human capital coefficient that cannot be corrected without outside information on the value of the parameter c . We will attempt to obtain one such estimate later on, but for now we will concentrate on the estimation of β , keeping in mind that this coefficient is likely to be a biased estimate of b , which is the parameter we are really interested in.

We will estimate a number of specifications based on (8) using different schooling series and a common set of other variables. Our first specification is obtained directly from (8) by replacing a_{it} by a set of period and country dummies,

$$(10) \quad q_{it} = \Gamma_1 + \gamma_i + \eta_{1t} + \alpha z_{it} + \beta h_{it} - \varphi e_{it} + \varepsilon_{1it}$$

where η_{1t} and γ_i are fixed time and country effects and ε_{1it} the disturbance term. A second specification is obtained by taking differences of (10),

$$(11) \quad \Delta q_{it} = \Gamma_2 + \eta_{2t} + \alpha \Delta z_{it} + \beta \Delta h_{it} - \varphi \Delta e_{it} + \lambda b_{it} + \varepsilon_{2it}$$

(which eliminates the country fixed effects), and a third one by extending (11) to allow for technological diffusion along the lines of de la Fuente (2002b). This last equation is of the form²⁷

$$(12) \quad \Delta q_{it} = \Gamma_3 + \mu_i + \eta_{3t} + \alpha \Delta z_{it} + \beta \Delta h_{it} - \varphi \Delta e_{it} + \lambda b_{it} + \varepsilon_{3it}$$

where Δ denotes annual growth rates (over the subperiod starting at time t) and b_{it} is a technological gap measure which enters the equation as a determinant of the rate of technical progress in order to allow for a catch-up effect. This term is the Hicks-neutral TFP gap between each country and the US at the beginning of each subperiod, given by

$$(13) \quad b_{it} = (q_{US,t} - \alpha z_{US,t} - \beta h_{US,t} + \varphi e_{US,t}) - (q_{it} - \alpha z_{it} - \beta h_{it} + \varphi e_{it}).$$

To estimate this specification we substitute (13) into (12) and use non-linear least squares with data on both factor stocks and their growth rates. In this specification the parameter λ measures the rate of (conditional) technological convergence. Notice that if this parameter is positive, relative TFP levels eventually stabilize, signalling a common asymptotic rate of technical progress for all countries, and

²⁷ Equation (12) is obtained from equation (11) under the assumption that the rate of technological progress (which is implicitly assumed to be constant across countries for each period in the first equation) is given by

$$\Delta a_{it} = \lambda b_{it} + \mu_i + \eta_{3t}$$

The term b_{it} is obtained by solving for TFP in the production function in log levels and taking differences with the US.

the country fixed effects μ_j capture permanent differences in relative total factor productivity that will presumably reflect differences in R&D investment and other omitted variables.

Our data on output, employment and investment are taken from an updated version of Dabán, Doménech and Molinas (1997), who replicate Summers and Heston's (1991) Penn World Table for the OECD using National Accounts data from this organization and a set of purchasing power parities specific to this sample.²⁸ The construction of the stock of physical capital using a perpetual inventory procedure is discussed in section 3 of Appendix 1. We use pooled data at five-year intervals starting in 1960 and ending either in 1985 or in 1990 depending on the duration of the different schooling series.

The results obtained with the different specifications are reported in Table 9. All equations contain period dummies and those in panels *b* and *d* of the table include fixed country effects as well. The equations shown in panel *a* correspond to equation (10) without fixed country effects and control for the employment ratio (e_{it}), which is highly significant and displays the expected negative sign. For the remaining equations, e_{it} turned out to be non-significant (which is not surprising given its very small time variation), so this variable is omitted (with very marginal changes in the remaining coefficients). Panel *e* contains results for the technological catch-up specification when only the significant country dummies are left in the equation so as to retain as much of the cross-country variation in the data as possible. This is our preferred specification. The last panel of the table shows average coefficient values and *t* ratios for each data set computed across the different specifications, and the last column within each panel reports average results across data sets for each specification.

The pattern of results that emerges as we change the source of the human capital data is consistent with our hypothesis about the importance of educational data quality for growth estimates. For all the data sets, the coefficient of human capital (β) is positive and significant in the specification in levels without fixed country effects (panel *a* of Table 9), but the size and significance of the estimates increases appreciably as we move to the data sets with higher reliability ratios. The differences are even sharper when the estimation is repeated with fixed country effects (panel *b* in Table 9) or with the data in growth rates with or without a catch-up effect (panels *c*, *d* and *e*). The results obtained with the Kyriacou, B&L and NSD data in growth rates are consistent with those reported by Kyriacou (1991), Benhabib and Spiegel (1994) and Pritchett (1999), who find insignificant (and sometimes negative) coefficients for human capital in an aggregate production function estimated with differenced data. On the other hand, our series and those of Cohen and Soto produce rather large and precise estimates of the human capital coefficient in most equations and, in the case of our preferred catch-up specification (panel *e*), yield plausible values of the remaining parameters of the model as well, with estimates of α close to the share of physical capital in national income and positive diffusion coefficients.

²⁸ The OECD has recently revised these series to incorporate changes in national accounting standards. As a result, our non-schooling data are slightly different from those used in de la Fuente and Doménech (2000 and 2001), and so are the results of the growth equations reported in this section.

Table 9: Alternative production function specifications**a. Log levels (without fixed country effects)**

<i>H data from:</i>	[a1]	[a2]	[a3]	[a4]	[a5]	[a6]	[a7]	[a8]	[a9]
	NSD	KYR	B&L93	B&L96	B&L00	C&S	D&D00	D&D02	avge.
α	0.580 (18.56)	0.588 (15.77)	0.512 (13.35)	0.512 (14.43)	0.479 (12.62)	0.447 (11.81)	0.451 (13.24)	0.448 (12.04)	0.502 (13.98)
β	0.078 (2.02)	0.186 (2.18)	0.141 (4.49)	0.165 (4.82)	0.238 (6.19)	0.397 (7.98)	0.407 (7.76)	0.378 (6.92)	0.249 (5.30)
φ	-0.257 (3.45)	-0.235 (3.07)	-0.311 (4.05)	-0.362 (5.19)	-0.432 (5.73)	-0.563 (7.46)	-0.560 (7.30)	-0.596 (6.97)	-0.414 (5.40)
adj. R^2	0.881	0.884	0.890	0.889	0.911	0.918	0.923	0.914	
std. error reg.	0.129	0.103	0.124	0.119	0.113	0.108	0.105	0.128	
no. of observ.	126	95	126	147	147	147	147	147	

- Note: based on equation (10) without fixed country effects.

b. Log levels with fixed country effects

<i>H data from:</i>	[b1]	[b2]	[b3]	[b4]	[b5]	[b6]	[b7]	[b8]	[b9]
	NSD	KYR	B&L93	B&L96	B&L00	C&S	D&D00	D&D02	avge.
α	0.540 (16.92)	0.533 (16.29)	0.531 (18.59)	0.539 (18.54)	0.536 (19.90)	0.516 (20.08)	0.506 (18.17)	0.491 (19.80)	0.524 (18.54)
β	0.068 (0.76)	0.066 (1.86)	0.136 (3.30)	0.115 (1.80)	0.203 (3.74)	0.608 (4.49)	0.627 (3.99)	0.958 (6.51)	0.348 (3.31)
adj. R^2	0.978	0.980	0.979	0.977	0.978	0.980	0.978	0.982	
std. error reg.	0.056	0.043	0.054	0.058	0.057	0.053	0.056	0.051	
no. of observ.	126	95	126	147	147	147	147	147	

- Note: based on equation (10) with a full set of fixed country effects.

c. Growth rates

<i>H data from:</i>	[c1]	[c2]	[c3]	[c4]	[c5]	[c6]	[c7]	[c8]	[c9]
	NSD	KYR	B&L93	B&L96	B&L00	C&S	D&D00	D&D02	avge.
α	0.504 (9.78)	0.534 (6.92)	0.501 (9.79)	0.490 (9.59)	0.493 (9.71)	0.477 (9.22)	0.475 (9.36)	0.470 (9.34)	0.493 (9.21)
β	0.079 (0.70)	0.009 (0.15)	0.089 (2.52)	0.083 (1.47)	0.079 (1.28)	0.525 (2.57)	0.520 (2.17)	0.744 (3.10)	0.266 (1.75)
adj. R^2	0.711	0.684	0.722	0.708	0.708	0.720	0.715	0.726	
std. error reg.	0.0097	0.0094	0.0095	0.0096	0.0096	0.0094	0.0094	0.0093	
no. of observ.	105	74	105	126	126	126	126	126	

- Note: based on equation (11).

Table 9: Alternative production function specifications (continued)**d. Growth rates with technological diffusion and fixed country effects**

<i>H data from:</i>	[d1]	[d2]	[d3]	[d4]	[d5]	[d6]	[d7]	[d8]	[d9]
	NSD	KYR	B&L93	B&L96	B&L00	C&S	D&D00	D&D02	avge.
α	0.406 (5.68)	0.397 (4.84)	0.421 (6.05)	0.330 (3.93)	0.327 (3.90)	0.414 (6.16)	0.399 (5.85)	0.400 (5.79)	0.387 (5.28)
β	-0.206 (1.61)	0.014 (0.29)	0.056 (1.80)	-0.007 (0.11)	-0.019 (0.31)	0.573 (3.52)	0.587 (3.47)	0.540 (2.89)	0.192 (1.24)
λ	0.094 (6.84)	0.144 (9.70)	0.096 (6.21)	0.069 (5.52)	0.068 (5.49)	0.097 (6.41)	0.091 (6.14)	0.094 (5.97)	0.094 (6.54)
adj. R^2	0.838	0.837	0.838	0.815	0.815	0.823	0.823	0.819	
std. error reg.	0.0073	0.0068	0.0073	0.0056	0.0076	0.0075	0.0074	0.0075	
no. of observ.	105	74	105	126	126	126	126	126	

- Note: based on equation (12) with a full set of fixed country effects.

e. Growth rates with technological diffusion and significant country dummies

<i>H data from:</i>	[e1]	[e2]	[e3]	[e4]	[e5]	[e6]	[e7]	[e8]	[e9]
	NSD	KYR	B&L93	B&L96	B&L00	C&S	D&D00	D&D02	avge.
α	0.435 (5.67)	0.433 (6.50)	0.490 (8.25)	0.357 (5.67)	0.360 (5.69)	0.372 (7.50)	0.343 (6.94)	0.345 (6.83)	0.392 (6.63)
β	-0.089 (1.39)	-0.008 (0.19)	0.062 (1.83)	-0.043 (0.63)	-0.039 (0.72)	0.300 (4.47)	0.399 (4.39)	0.394 (4.57)	0.122 (1.54)
λ	0.083 (8.98)	0.144 (7.36)	0.091 (7.40)	0.065 (10.2)	0.065 (10.2)	0.084 (7.16)	0.075 (6.72)	0.074 (7.07)	0.085 (8.14)
adj. R^2	0.840	0.844	0.842	0.820	0.820	0.823	0.823	0.828	
std. error reg.	0.0072	0.0066	0.0072	0.0075	0.0075	0.0074	0.0074	0.0073	
no. of observ.	105	74	105	126	126	126	126	126	

- Note: based on equation (12). Only significant country dummies are left in the equation.

f. Averages across specifications

<i>H data from:</i>	[f1]	[f2]	[f3]	[f4]	[f5]	[f6]	[f7]	[f8]
	NSD	KYR	B&L93	B&L96	B&L00	C&S	D&D00	D&D02
β	-0.014	0.053	0.097	0.063	0.092	0.481	0.508	0.603
(t)	(0.10)	(0.80)	(2.79)	(1.47)	(2.04)	(4.61)	(4.36)	(4.80)

Notes:

- All equations include period dummies.

- White's heteroscedasticity-consistent t ratios in parentheses below each coefficient.


- The average value of t shown in panel f is computed respecting the sign of the t ratios obtained for the different specifications; i.e. for this computation we assign to each t ratio the same sign as the corresponding coefficient estimate.

- In panel a, ϕ is the coefficient of the ratio of employment to the population aged 15 to 64 (e).

- Key: NSD = Nehru et al (1995); KYR = Kyriacou (1991); B&L = Barro and Lee (various years); C&S = Cohen and Soto (2001); D&D = de la Fuente and Doménech (various years); D&D02 refers to this paper.

We have checked the robustness of the results obtained with our latest data set in a number of ways. First, we have reestimated our preferred specification (the catch-up equation labeled [e8] in Table 9e) for all the possible subsamples obtained by deleting one country at a time from the original data set. Figure 9 displays the estimated human capital coefficient and the 95% confidence interval around it with the coefficient estimates arranged in decreasing order across subsamples. As can be seen in the Figure, although there are two particularly influential observations (Greece and Portugal), our estimate of β remains significantly different from zero at conventional confidence levels even in the most unfavourable case. By contrast, Temple (1998) reports that Mankiw, Romer and Weil's (1992) proxy for educational investment loses its significance once a few influential observations are removed. He shows, in particular, that the removal of Japan from the OECD sample is sufficient to make the coefficient of the human capital variable insignificant (with a t ratio below one).

Figure 9: Estimated coefficient of human capital and 95% confidence interval around it when deleting one country at a time from the sample



- Note: Catch-up specification with country dummies (equation [e8] in Table 9e), estimated after excluding the country shown in the horizontal axis. Only significant country dummies are left in each equation.
- Legend: Gr = Greece; Sw = Sweden; Ir = Ireland; No = Norway; Ge = West Germany; Ost = Austria; Be = Belgium; NI = Netherlands; US = United States; Fr = France; Fi = Finland; CH = Switzerland; NZ = New Zealand; Aus = Australia; It = Italy; UK = United Kingdom; Dk = Denmark; Ja = Japan; Ca = Canada; Sp = Spain; Po = Portugal.

Second, we have replicated the exercise using the Mincerian specification that seems to be favoured by many authors in the recent literature. This specification change, which involves replacing the log of H by its level in equations (10)-(13), produces results that are qualitatively very similar to those discussed above (and slightly worse all around). (See Table A.8 in section 1 of Appendix 2). Finally, we have reestimated our preferred model instrumenting the growth rate of years of schooling with the (log) stock of the same variable at the beginning of the current subperiod without any

appreciable change in the estimated coefficients of the model (see Table A.9 in Appendix 2). We interpret this finding as an indication that our results are not primarily driven by reverse causation bias.²⁹

5. Correcting for measurement error bias

The results summarized in Table 9 strongly suggest that measurement error often induces a large downward bias in human capital coefficients. They also show that improvements in data quality reduce this bias and generate results that are generally more favourable to the view that investment in schooling contributes substantially to productivity growth. To make this point visually, Figure 10 plots the various estimates of β given in panels *a* through *d* of Table 9 against the relevant SUR reliability ratios, along with the regression lines fitted for each of the growth specifications.³⁰ The scatter shows a clear positive correlation between these two variables within each specification and suggests that the true value of β is at least 0.50 (which is the prediction of the levels equation for $r = 1$).

In the remainder of this section we will explore more systematically the relationship between data quality and growth equation results. We will use an extension of the classical errors-in-variables model (CEVM) to obtain a set of meta-estimates of β after correcting for measurement error bias. To proceed further we will drop the catch-up specification (panel *d* of Table 9). There are two reasons for this choice. The first one is that, as suggested by Figure 10, an extrapolation based on our estimates of β with this specification and the relevant reliability ratios will yield an extremely high and rather implausible meta-estimate of β . The second and more important reason is that the catch-up model is a restricted non-linear specification in which human capital enters both in growth rates and in levels (through the technological gap variable). Hence, we would have to deal with two sources of measurement error, and with features (non-linearity and parameter restrictions) that are not easily incorporated into the errors-in-variables model we will be using below.

We will write the model we want to estimate (i.e. the different versions of the growth equation) in the generic form

$$(14) \quad Q = H\beta + \mathbf{X}\alpha + u_1$$


where Q is (some transformation of) output per employed worker, H the true stock of human capital, $\mathbf{X} = (Z, E)$ a row vector whose components are the other growth regressors (the stock of capital per

²⁹ While this is a standard concern in this sort of exercise, we believe it is very unlikely that it will induce a significant upward bias in our preferred specification. There are two reasons for this: one is the relatively high frequency at which our observations are taken, and the other is that our estimates are obtained using a fairly complete specification of the technical progress function that allows for technological diffusion and for country fixed effects that should help control for omitted variables such as R&D investment. The first of these features makes it unlikely that any increases in enrollment induced by higher income or faster growth will have time to appreciably affect schooling stocks. The second feature is important because the reverse causation problem arises when the residual of the growth equation is not a "clean" random disturbance but contains systematic components of the growth rate that will enter the demand for education because they can be anticipated by individuals. See de la Fuente (2002a) for a more detailed discussion of these issues.

³⁰ We do not use the results in block *e* of Table 9 because each of these equations contains a different set of country dummies. This makes the coefficients not strictly comparable with each other and raises some questions about the appropriate reliability ratio.

employed worker, Z , and the employment ratio, E) and α a column vector of coefficients. It will be assumed that the error term u_1 satisfies all the standard assumptions of the linear regression model (and is, in particular, uncorrelated with the regressors) so that the estimation of (14) by OLS with the correctly measured stock of human capital will be consistent.

Figure 10: Estimated β vs. SUR reliability ratio



We want to calculate the bias that arises when β is estimated in equation (14) using a noisy proxy for the stock of human capital,

$$P_j = H + \varepsilon_j$$

rather than H itself. We will assume that the measurement error terms, ε_j , have the following structure:

$$(15) \varepsilon_j = \omega_j + \rho_j \varepsilon + \mathbf{X} \delta_j$$

where ω_j is an idiosyncratic error component and ρ_j a coefficient that measures the extent to which data set j amplifies or dampens a common source of error (say, that present in the underlying primary data) which is captured by an iid disturbance, ε . Notice that this specification allows measurement error to be correlated across data sets (as $E \varepsilon_j \varepsilon_k$ will be different from zero if $\rho_j \rho_k \neq 0$) and with the components of $\mathbf{X} = (Z, E)$, as indicated by the last term of (15) where $\delta_j = (\delta_{jz}, \delta_{je})'$ is a column vector

of coefficients. Finally, it will be assumed that both the common and the idiosyncratic components of the measurement error terms are uncorrelated with each other and with H and \mathbf{X} , i.e. that

$$(16) E H \varepsilon = E H \omega_j = E \omega_j \varepsilon = E \omega_j \omega_k = E X_n \varepsilon = E X_n \omega_j = 0$$

for all j and $k \neq j$ and for all components X_n of \mathbf{X} .

Under these assumptions, it is shown in section 2 of Appendix 2 that the expected value in large samples (probability limit) of the OLS estimator of β obtained using schooling series P_j is given by³¹

$$(17) \text{plim } \hat{\beta}_j = \frac{(1 - ER_H^2)r_j'}{1 - ER_H^2 r_j'} \beta \equiv a_j \beta$$

where β is the true value of the parameter, $ER_H^2 \equiv \text{plim } R^2(H|\mathbf{X})$ is the probability limit of the coefficient of determination of a regression of H on the rest of the explanatory variables of the growth equation, \mathbf{X} , and r_j' is what we will call the *adjusted reliability ratio* of series P_j . This variable is defined as the standard reliability ratio of the series $P_j' = P_j - \mathbf{X}\delta_j$ that is obtained by removing from the original schooling series P_j the component of measurement error that is correlated with \mathbf{X} ; that is,

$$(18) r_j' = \frac{EH^2}{E(P_j - \mathbf{X}\delta_j)'(P_j - \mathbf{X}\delta_j)}$$

We will refer to the expression that multiplies the true value of the parameter in the right-hand side of (17) as the *attenuation coefficient*, and we will denote it by a_j . Notice that the attenuation coefficient reduces to the reliability ratio when P_j is the only regressor in the growth equation (or when H is perfectly uncorrelated with the rest of the right-hand side variables). Otherwise, there is a further loss of signal that increases with the correlation of H with the other explanatory variables in the model. Using the fact that ER_H^2 and r_j' must both lie between zero and one, it is easily seen that a_j is an increasing function of r_j' , with $0 \leq a_j \leq r_j' \leq 1$, and that $a_j < r_j'$ except if $r_j' = 0$, $r_j' = 1$ or $ER_H^2 = 0$. Hence, contrary to what may appear to be the case in equation (17') in footnote 31, measurement error cannot reverse coefficient signs in large samples, even when the human capital proxy is highly correlated with other regressors.

Our meta-estimates of β will be obtained by estimating (for each of the different specifications of the production function we have used) a regression of the form

$$(19) \hat{\beta}_j = \beta \tilde{a}_j + \eta_j$$

where η_j is a disturbance term and \tilde{a}_j a consistent estimate of the relevant attenuation coefficient which will be constructed independently of $\hat{\beta}_j$. This equation is the finite sample counterpart of equation (17), and it is obtained from this expression by deleting the probability limit, leaving on the left-hand side the corresponding sample estimate, and by replacing the attenuation coefficient on the right-hand side by a consistent estimate of it and adding a disturbance term to capture small-sample

³¹ When measurement error is uncorrelated with the components of \mathbf{X} (i.e. when $\delta = \mathbf{0}$) equation (17) can be shown to be equivalent to the bias-correction formula given in Krueger and Lindhal (2001) which, in our notation, would be written

$$(17') \text{plim } \hat{\beta}_j = \frac{r_j - ER_j^2}{1 - ER_j^2} \beta$$

deviations from the asymptotic relationship between these two variables. The following sections discuss the construction of \tilde{a}_j . Technical details can be found in section 2 of Appendix 2.

5.1. Estimating reliability ratios with correlated errors

We will first discuss the estimation of the unadjusted reliability ratios. When measurement error in H is correlated across data sets or with the other regressors of the growth equation, the procedure developed by Krueger and Lindhal (2001) will generally yield asymptotically biased estimates of reliability ratios. Consistent estimates of these ratios, however, can be recovered through a two-step procedure. We will first obtain all possible pairwise estimates of the reliability ratios of the different schooling data sets by regressing them on each other as suggested by Krueger and Lindhal. As part of the first stage, we will also estimate a set of auxiliary regressions that will exploit the information contained in the correlations between the different schooling series and the remaining growth regressors. The coefficient estimates obtained in this manner will be biased, but they can be used as data in a set of second-stage regressions that will yield consistent estimates of the parameters of interest. As equation (19), these second-stage equations are derived from the probability limits of the first-stage estimators.

As noted, the first step in the estimation procedure involves regressing the different schooling series on each other and on the remaining explanatory variables of the growth model. We fix some data set P_j and use it to try to explain the remaining data sets $k \neq j$,³² as well as the other growth regressors contained in the vector \mathbf{X} . Hence, for each j we estimate by OLS the following set of equations:

$$(20) P_k = r_{jk} P_j + u_{jk} \quad \text{for } k = 1, \dots, J \text{ with } k \neq j \text{ and}$$

$$(21) X_n = \mu_{jn} P_j + u_{jn} \quad \text{for } n = z, e$$

where the u 's are disturbance terms and J ($= 8$) the number of alternative proxies for H that are available. This yields (inconsistent) estimates of r_j and μ_n that we will denote by \hat{r}_{jk} and $\hat{\mu}_{jn}$ (hats will be used throughout to indicate first-stage OLS estimates).³³ In addition to the J systems of the form given in (20)-(21), we also estimate by OLS all the "reverse" regressions of P_j on \mathbf{X} ,

$$(22) P_j = \mathbf{X}\phi_j + u_{Xj}$$

to obtain coefficient estimates we will denote by $\hat{\phi}_j$.

It is easily shown that

$$(23) \text{plim } \hat{\phi}_j = \phi + \delta_j$$

where ϕ is the probability limit of the vector of coefficients of \mathbf{X} in a regression analogous to (22) in which P_j is replaced by H . Hence, the estimation of equation (22) yields consistent estimates of δ_j "up to a constant" -- that is, differences across series in the estimated coefficients $\hat{\phi}_j$ will be consistent

³² Contrary to what we did in Section 3.3, we now construct all possible pairwise estimates of the reliability ratio, without excluding series that are members of the same family of data sets.

³³ As elsewhere in this paper, these equations include period dummies and, where appropriate, country dummies as well in order to remove time and/or country means.

estimates of differences in the value of δ_j . To exploit this information, we take a specific data set (D&D02) as a reference, denote the corresponding value of δ_j by δ , define Δ_j by

$$(24) \Delta_j = \delta_j - \delta$$

and obtain a consistent estimate of its value as

$$(25) \tilde{\Delta}_j = \hat{\phi}_j - \hat{\phi}_{DD02}.$$

This leaves us with only the two components of $\delta = (\delta_z, \delta_e)'$ to be estimated in the second stage.

To derive the second-stage regressions, we begin by calculating the probability limits of the first-stage estimators, which are given by (see Appendix 2):

$$(26) \text{plim } \hat{r}_{jk} = r_j [1 + e_j e_k + (\delta_{\mathbf{k}} + \delta_j)' \mu] + C_{jk}$$

$$(27) \text{plim } \hat{\mu}_{jn} = r_j \mu_n + C_{jn}$$

where $\mu = (\mu_z, \mu_e)'$ and r_j are the true values of the parameters of interest and the remaining terms in these expressions are given by

$$(28) C_{jk} \equiv \frac{\delta_j' (E\mathbf{X}'\mathbf{X}) \delta_{\mathbf{k}}}{EP_j^2} \quad C_{jn} \equiv \frac{\delta_j' E\mathbf{X}' X_n}{EP_j^2} \quad \text{and} \quad e_j \equiv \rho_j \sqrt{\frac{E\varepsilon^2}{EH^2}}$$

Hence, the pairwise estimate of the reliability ratio of series P_j obtained using P_k as a reference, \hat{r}_{jk} , will be biased upward (downward) whenever the measurement error components of the two series are positively (negatively) correlated, and an additional bias will be introduced by the existence of a non-zero correlation of either series with the components of \mathbf{X} .

Equations (26) and (27) relate the expected values of \hat{r}_{jk} and $\hat{\mu}_{jn}$ to the true values of r_j and μ_n and to the coefficients that describe the structure of the measurement error terms, e_j and $\delta_j = \delta + \Delta_j$. Proceeding as above, we construct the second-stage regressions as the natural finite sample approximations to these relations; that is, we suppress the probability limits in the left-hand side of (26) and (27), leaving as dependent variables the actual first-stage parameter estimates, replace any population moments that appear on the right-hand side by their sample counterparts, and introduce a disturbance term. Rewriting C_{jk} and C_{jn} as explicit functions of δ and Δ_j , the equations to be estimated can be written

$$(29) \hat{r}_{jk} = r_j [1 + e_j e_k + 2(\tilde{\delta}_z \mu_z + \tilde{\delta}_e \mu_e) + (\tilde{\Delta}_{jz} + \tilde{\Delta}_{kz}) \mu_z + (\tilde{\Delta}_{je} + \tilde{\Delta}_{ke}) \mu_e] + \frac{1}{\text{svar } P_j} \left\{ \tilde{d}_{jk} + \tilde{V}_{zz} \delta_z^2 + 2\tilde{V}_{ze} \delta_z \delta_e + \tilde{V}_{ee} \delta_e^2 + (\tilde{d}_{jz} + \tilde{d}_{kz}) \delta_z + (\tilde{d}_{je} + \tilde{d}_{ke}) \delta_e \right\} + \eta_{jk}$$

$$(30) \hat{\mu}_{jn} = r_j \mu_n + \frac{1}{\text{svar } P_j} \left\{ \tilde{d}_{jn} + \tilde{V}_{nz} \delta_z + \tilde{V}_{ne} \delta_e \right\} + \eta_{jn}$$

where η_{jk} and η_{jn} are disturbance terms and tildes denote consistent sample estimates of the relevant quantities. Notice that we have replaced EP_j^2 by the sample variance of P_j (denoted by $\text{svar } P_j$) and $\mathbf{V} \equiv E\mathbf{X}'\mathbf{X}$ by the sample variance-covariance matrix of $\mathbf{X} = (Z, E)$. This matrix will be written

$$\tilde{\mathbf{V}} = \begin{pmatrix} \tilde{V}_{zz} & \tilde{V}_{ze} \\ \tilde{V}_{ze} & \tilde{V}_{ee} \end{pmatrix} = (\tilde{\mathbf{V}}_z, \tilde{\mathbf{V}}_e)$$

where $\tilde{\mathbf{V}}_n$ is its n -th column. Finally, the terms \tilde{d}_{jk} and \tilde{d}_{jn} are defined as

$$(31) \tilde{d}_{jk} \equiv \tilde{\Delta}_j' \tilde{\mathbf{V}} \tilde{\Delta}_k \quad \text{and} \quad \tilde{d}_{jn} \equiv \tilde{\Delta}_j' \tilde{\mathbf{V}}_n$$

and can be computed using information known at this stage.

Thus, the only unknown quantities in (29) and (30) are the coefficients to be estimated: r_j and e_j for $j = 1 \dots J$, μ_z , μ_e , δ_z and δ_e . We estimate both equations jointly by "stacking them" so that, for each j , the first J observations of the dependent variable (one of which will be missing as k must be different from j) correspond to the pairwise estimates of the reliability ratio, \hat{r}_{jk} , and the last two observations correspond to the first-stage estimates, $\hat{\mu}_{jn}$. Next, we recover an estimate of ϕ using (23) as

$$(32) \quad \tilde{\phi} = \hat{\phi}_{DD02} - \tilde{\delta}.$$

We also estimate a restricted version of (29) and (30) in which we impose the assumption that measurement error is uncorrelated with \mathbf{X} (i.e. that $\delta_j = \mathbf{0}$ for all j). When this is done, a similar restriction is imposed on equation (22) to obtain a direct estimate of $\tilde{\phi}$. The detailed results of the estimation of the restricted and unrestricted models for each growth specification can be found in section 2.f of Appendix 2.

5.2 Adjusted reliability ratios and attenuation coefficients

We will refer to the reliability ratio estimates obtained from equations (29) and (30) as the (*restricted or unrestricted*) "consistent" estimates of the reliability ratio, and we will denote them by \tilde{r}_j . One drawback of these estimates is that they depend on the correlation of each schooling series with \mathbf{X} . As a result, they do not necessarily lie between zero and one when δ is different from zero, and are not strictly comparable either across data sets or with alternative estimates of the same magnitude that assume $\delta = \mathbf{0}$. To get around these problems, it is convenient to work with the adjusted reliability ratios we have defined at the beginning of this section, r_j' . Once these adjusted ratios have been computed, the attenuation coefficient that appears in equation (17) can be calculated. We show in Appendix 2 that these variables can be estimated by

$$(33) \quad \tilde{r}_j' = \frac{\tilde{r}_j}{(1 - \tilde{C}_{jj}) - \tilde{\mathcal{X}}_j' \tilde{\mu} \tilde{r}_j}$$

and

$$(34) \quad ER_H^2 = \tilde{\mu}' \tilde{\phi}$$

where \tilde{C}_{jj} is constructed using (28), our estimates of $\delta_j = \delta + \Delta_j$ and the sample variances and covariances of P_j and \mathbf{X} .

Table 10 collects our different estimates of the reliability ratios and attenuation coefficients of the different schooling series used in the literature. These variables are shown in relative terms, with the average value of each series across data sets (shown at the bottom of the table under *avge. value*) normalized to 100. The three panels of the table refer to the different growth specifications or data transformations we have retained in this section (levels, fixed effects and differences (growth rates)). Within each panel, the left-hand side block of the table shows three alternative estimates of the adjusted reliability ratio: the SUR estimates constructed in section 3.3, and the two "consistent" estimates obtained using the model developed in this section. The "restricted" consistent estimate

assumes that measurement error is uncorrelated with the remaining growth regressors ($\delta_j = \mathbf{0}$ for all j), whereas the unrestricted estimate allows these parameters to be different from zero. For the unrestricted estimate, we show the adjusted reliability ratio computed as indicated above. In the other two cases, the raw and adjusted reliability ratios are the same.

The right-hand side block of each panel of the table shows the attenuation coefficient (a_j) implied by each of the three estimates of the reliability ratio we have just discussed. The value of a_j is calculated using equation (17) and an estimate of ER_H^2 . For the case of the consistent reliability ratios, ER_H^2 is generally obtained as $\tilde{\mu}'\tilde{\phi}$ using equation (34) but there are some exceptions, as in some cases this procedure yields estimates of ER_H^2 that do not lie between zero and one (see the notes to the table). In the case of the SUR reliability ratios, the value of ER_H^2 cannot be obtained in the same way because no estimate of μ is then available. Under the assumptions used to estimate these ratios, however, it can be shown that (see equation (26) in Appendix 2)

$$(35) ER_j^2 = r_j ER_H^2$$

so we obtain an estimate of ER_H^2 by regressing the observed value of $R_j^2 = R^2(P_j|\mathbf{X})$ on the SUR estimate of r_j without a constant. The last two rows of each panel show, respectively, the value of $\tilde{\mu}'\tilde{\phi}$ and the value of ER_H^2 that was used to construct the displayed attenuation coefficients.

At the bottom of each panel of the table we show the average value (across data sets) of the reliability ratio and the attenuation coefficient, and the correlation between each of these series and the SUR reliability ratio. These statistics show that, when the data are used in differences or after removing fixed effects, the corrections required for the presence of additional growth regressors and for correlated measurement errors are generally minor. Hence, the SUR reliability ratios are decent measures of the information content of the different series for these data transformations. For the data in levels, however, the situation is more complicated. Our restricted consistent model suggests that the SUR reliability ratios are highly inflated by the positive correlation of measurement error across data sets (notice that the average value of the reliability ratio drops by two thirds as we go from the first to the second column of panel a of the table). When we allow the measurement error term to be correlated with \mathbf{X} , however, the estimated values of the (adjusted) reliability ratio rise above our SUR estimates, even though there are still clear signs of positively correlated errors (see Table A.11 in Appendix 2). We find the size of the correction for the correlation of measurement error with \mathbf{X} surprising, particularly because the assumption that $\delta_j = \mathbf{0}$ seems fairly reasonable ex-ante, but there are clear indications that this assumption does not hold for the data in levels, particularly in regard to the employment ratio (see Tables A.10a and A.11 in Appendix 2).

On the whole, our alternative estimates of reliability ratios display a fairly consistent picture of the relative quality of the different series and reinforce our earlier conclusion that recent attempts to increase the information content of the schooling data have been fairly successful. There are, however, some patterns in the results that suggest that our assumptions about the nature of measurement error

Table 10: Adjusted reliability ratios and attenuation coefficients, relative values of various estimates

a. data in levels

	<u>reliability ratios</u>			<u>attenuation coefficients</u>		
	<i>SUR</i>	consistent restricted	consistent unrestr.	<i>SUR</i>	consistent restricted	consistent unrestr.
NSD	53.6	84.0	85.2	38.6	80.4	53.6
KYR	97.3	171.3	128.5	88.1	186.0	229.2
B&L93	70.7	52.6	75.6	55.4	48.2	41.6
B&L96	83.6	64.4	82.1	70.1	60.1	49.3
B&L00	97.7	70.9	89.3	88.7	66.7	59.8
C&S	147.9	122.6	113.3	190.5	123.9	122.6
D&D00	123.4	118.9	110.6	132.0	119.5	112.1
D&D02	125.8	115.3	115.3	136.8	115.2	132.0
avge. value	0.617	0.221	0.741	0.421	0.116	0.302
corr. w/ <i>SUR</i>	1.000	0.526	0.696	0.984	0.498	0.479
$\tilde{\mu}'\tilde{\phi}$					1.069	0.887
ER_H^2				0.607	0.551	0.887

- Note: Since the value of $\tilde{\mu}'\tilde{\phi}$ exceeds one in the restricted consistent case, we use as ER_H^2 the highest observed value of R_j^2 . Under the assumptions used in this case, equation (35) holds so we should have $ER_j^2 < ER_H^2$ for any P_j with $r_j < 1$ and this implies that the reported a_j 's should overestimate the true ones. We do not use equation (35) to estimate ER_H^2 because this procedure also yields a value of this parameter greater than one.

Table 10: Adjusted reliability ratios and attenuation coefficients (continued)

b. fixed effects

	<u>reliability ratios</u>			<u>attenuation coefficients</u>		
	<i>SUR</i>	consistent restricted	consistent unrestr.	<i>SUR</i>	consistent restricted	consistent unrestr.
NSD	25.9	28.1	46.0	22.9	25.4	46.0
KYR	30.2	47.2	16.0	26.8	43.2	16.0
B&L93	33.1	22.7	17.5	29.4	20.5	17.5
B&L96	39.9	53.1	41.4	35.7	48.8	41.4
B&L00	48.5	57.5	47.9	43.6	53.0	47.9
C&S	149.2	141.8	175.2	145.3	139.0	175.2
D&D00	223.0	219.4	144.6	231.0	228.4	144.6
D&D02	250.2	230.2	311.5	265.3	241.7	311.5
avge. value	0.367	0.416	0.301	0.339	0.393	0.301
corr. w/ <i>SUR</i>	1.000	0.994	0.921	0.999	0.995	0.921
$\tilde{\mu}'\tilde{\phi}$					0.163	-0.020
ER_H^2				0.198	0.163	0.000

- Note: Since the value of $\tilde{\mu}'\tilde{\phi}$ is negative in the unrestricted consistent case, we use a zero value for ER_H^2 .

c. differences

	reliability ratios			attenuation coefficients		
	<i>SUR</i>	consistent <i>restricted</i>	consistent <i>unrestr.</i>	<i>SUR</i>	consistent <i>restricted</i>	consistent <i>unrestr.</i>
NSD	26.0	26.1	48.3	25.2	25.6	48.1
KYR	25.5	17.1	9.9	24.8	16.7	9.8
B&L93	11.5	17.2	15.4	11.2	16.8	15.3
B&L96	15.5	45.4	41.2	15.0	44.6	41.0
B&L00	17.6	53.7	52.8	17.1	52.8	52.6
C&S	184.1	220.6	221.3	183.0	221.7	221.5
D&D00	216.8	172.5	150.1	216.7	172.2	150.0
D&D02	303.0	247.6	261.1	307.0	249.7	261.7
avge. value	0.254	0.330	0.301	0.247	0.325	0.300
corr. w/ <i>SUR</i>	1.000	0.962	0.951	1.000	0.962	0.951
$\tilde{\mu}'\tilde{\phi}$					0.039	0.009
ER_H^2				0.060	0.039	0.009

- Note: In all panels of the table, reliability ratios and attenuation coefficients are normalized in each column by their average value taken across data sets.

may not be entirely adequate for the estimation of reliability ratios. Perhaps the clearest indication of this can be found in the behaviour of the relative reliability ratios for the Kyriacou and NSD series in levels, which rise sharply as we go from the *SUR* to the consistent estimates -- to the extent that, as we will see below, the Kyriacou series becomes an extreme outlier when we consider the relationship between consistent attenuation coefficients and human capital coefficients for the specification in levels. We suspect that this anomaly may be due to our assumption of a common source of error arising from reliance on similar primary data (see equation (15) above). Since the Kyriacou and NSD series, unlike the rest of the data sets, rely mostly on enrollment data rather than on census compilations, their estimated correlation with the primary error shared by the remaining data sets (captured by the parameter ρ_j) is likely to be quite low. This, in turn, will cause the consistent reliability ratio estimates for these two series to increase relative to those of the remaining data sets (see equation (16)).³⁴

5.3 Meta-estimates of β

Table 11 shows the meta-estimates of β obtained through the estimation of equation

$$(19) \hat{\beta}_j = \beta \tilde{a}_j + \eta_j$$

which is reproduced here for convenience. The table has three panels that correspond to the alternative estimates of the attenuation coefficient shown in Table 10. Within each panel, the columns correspond to the three growth specifications we have been working with (levels, fixed effects and differences). In

³⁴ An additional peculiarity of the Kyriacou data set is that it is the only one that refers to the attainment of the labour force, rather than that of the adult population. Under our assumptions, this is likely to translate into a high negative correlation of measurement error with the employment ratio.

panels *b* and *c*, we add an additional levels specification that includes a dummy for the Kyriacou data set which, as noted in the previous section, is an extreme outlier for the data in levels. This fact is also illustrated in Figure 11, where we show the scatter diagram obtained with the unrestricted consistent attenuation coefficients together with two alternative regression lines fitted to the data in levels with and without the Kyriacou observation.

Table 11: Meta-estimates of β

a. with SUR attenuation coefficients

	[1] levels	[2] f. e.	[3] diffs.	[4] levels	[5] f. e.	[6] diffs.
β	0.569 (6.30)	0.966 (10.05)	0.965 (16.31)	0.587 (16.50)	0.998 (16.12)	1.016 (22.23)
constant	0.009 (0.22)	0.020 (0.46)	0.028 (1.28)			
R^2	0.869	0.944	0.978	0.868	0.942	0.972

b. with restricted consistent attenuation coefficients


	[1] levels	[2] levels	[3] f. e.	[4] diffs.	[5] levels	[6] levels	[7] f. e.	[8] diffs.
β	0.954 (1.03)	3.224 (3.77)	0.898 (7.76)	0.872 (12.52)	1.967 (5.31)	2.606 (10.05)	0.889 (12.80)	0.843 (18.82)
constant	0.138 (1.19)	-0.070 (0.76)	-0.001 (0.09)	-0.018 (0.58)				
dummy Kyr		-0.438 (3.46)				-0.375 (4.07)		
R^2	0.151	0.749	0.909	0.963	-0.050	0.721	0.909	0.961

c. with unrestricted consistent attenuation coefficients

	[1] levels	[2] levels	[3] f. e.	[4] diffs.	[5] levels	[6] levels	[7] f. e.	[8] diffs.
β	0.238 (0.93)	1.078 (5.58)	1.054 (10.62)	0.940 (10.51)	0.672 (4.33)	1.053 (14.77)	1.107 (16.55)	0.912 (16.06)
constant	0.177 (1.97)	-0.007 (0.14)	0.031 (0.74)	-0.016 (0.43)				
dummy Kyr		-0.551 (5.28)				-0.541 (7.66)		
R^2	0.125	0.867	0.949	0.948	-0.438	0.867	0.945	0.947

We first estimate equation (19) with a constant, which is generally not significant as predicted by the model, and then without the constant to obtain our final estimates of the coefficient of human capital. These meta-estimates of β consistently display very large values and significantly exceed those found in the previous literature. If we leave aside the atypical results obtained with the data in levels and the consistent attenuation coefficients when the Kyriacou outlier is included in the sample, our estimates of β range from 0.587 to 1.967, and from 0.843 to 1.107 if we restrict ourselves to the results obtained with the fixed effects and the differenced growth models, which are less likely to be misspecified and do not display obvious outliers.

Figure 11: Estimated β vs. unrestricted consistent attenuation coefficient



Before commenting further on these results, it should be recalled that our use of data on population attainment (rather than the attainment of employed workers) potentially introduces a bias. As noted in section 4, the production function parameter we would like to estimate is not β itself, but rather $b = \beta/c$ where c is one of the coefficients of a regression of the form

$$(7) \text{ } he_{it} = c \cdot h_{it} - d \cdot e_{it}$$

where h is the log of population attainment, he the log of the average attainment of employed workers and e the ratio of employment to the adult population. Since the data sets we have used do not provide information on the schooling of employed workers, we have used panel data for the Spanish regions from Mas et al (1998) to estimate equation (7) thus obtaining an outside estimate of c that may give us

an idea of the size of the required correction. Since measurement error is almost certainly a problem in these data as well (and there are no alternative sources that can be used to estimate reliability ratios), we obtain a set of bounds on the value of c by estimating both (7) and the reverse regression with the roles of h and h_e interchanged. The exercise is repeated with the data in (log) levels with and without fixed effects and in growth rates to obtain three alternative intervals of plausible values of c that can be matched with our three growth specifications to obtain ranges of possible values of b .

Table 12: range of meta-estimates of b

	<i>SUR</i>			<i>restricted</i>			<i>unrestricted</i>		
	<i>levels</i>	<i>f. e.</i>	<i>diffs.</i>	<i>levels</i>	<i>f. e.</i>	<i>diffs.</i>	<i>levels</i>	<i>f. e.</i>	<i>diffs.</i>
\hat{c} , <i>direct</i>	0.907	0.867	0.983	0.907	0.867	0.983	0.907	0.867	0.983
\hat{c} , <i>reverse</i>	0.993	1.182	1.468	0.993	1.182	1.468	0.993	1.182	1.468
$\tilde{\beta}$	0.587	0.998	1.016	2.606	0.889	0.843	1.053	1.107	0.912
b_{min}	<i>0.591</i>	<i>0.844</i>	<i>0.692</i>	<i>2.624</i>	<i>0.752</i>	<i>0.574</i>	<i>1.060</i>	<i>0.937</i>	<i>0.621</i>
b_{max}	0.647	1.151	1.034	2.875	1.025	0.858	1.162	1.277	0.928

The results are shown in Table 12. The first two rows of the table show the bounds on the value of c obtained from the direct and reverse regressions based on (7). The third row shows the estimates of β taken from Table 11, and the last two rows display the minimum and maximum values of b , which are obtained by dividing the estimated β by each of the estimates of c . As can be seen in the table, the correction is potentially important in some cases but does not qualitatively change our previous results. Our smallest lower bound for b (0.574) is roughly twice as large as Mankiw, Romer and Weil's (1992) estimate of $1/3$, which could probably have been considered a consensus value for this parameter a few years ago and has lately come to be seen as too optimistic in the light of recent negative results in the literature.

We interpret our results as a clear indication that investment in human capital is an important growth factor whose effect on productivity has been underestimated in previous studies because of poor data quality. Considerable uncertainty remains, however, regarding the correct value of the coefficient of this factor in an aggregate production function. Values of b of the order of 0.6 imply Mincerian rates of return on schooling (of around 6%) that are consistent with the evidence available from microeconomic wage equations.³⁵ Somewhat higher values of this parameter remain plausible as our macroeconomic estimates should pick up any externalities that may be associated with the accumulation of human capital, such as those arising from the contribution of an educated labour force to the development and adoption of new technologies. But some of the estimates at the upper end of our range of results seem decidedly implausible. Possible explanations for these very large coefficient values include the underestimation of the relevant reliability ratios (for the restricted model in levels)

³⁵ See for instance Card (1999) and Harmon, Walker and Westergaard-Nielsen (2001).

and an upward bias arising from reverse causation. An additional problem we have not addressed in this paper is that the likely mismeasurement of the remaining growth regressors may also bias the coefficient of human capital.

6. Conclusion

Existing data on educational attainment contain a considerable amount of noise. Due to changes in classification criteria and other inconsistencies in the primary data, the most widely used schooling series often display implausible time-series and cross-section profiles. After discussing the methodology and contents of the most widely used schooling data sets and documenting some of their weaknesses, we have constructed new estimates of educational attainment for a sample of OECD countries. We have attempted to increase the signal-to-noise ratio in these data by exploiting a variety of sources not used by previous authors, and by eliminating to the extent that it was possible sharp breaks in the series that must reflect changes in data collection criteria. While our estimates unavoidably involve a fair amount of guesswork, we believe that they provide a more reliable picture of cross-country relative educational attainments and of their evolution over time than previously available data sets.

We have also constructed statistical measures of the information content of the schooling data sets used in the growth literature under alternative assumptions on the nature of measurement error. While the choice of assumptions does make some difference for the ranking of the different data sets, on the whole these indices support our view that the amount of measurement error in these data is rather large, and clearly suggest that both our attainment series and those constructed by Cohen and Soto (2001) constitute a significant improvement over earlier sources. Even so, remaining discrepancies among recent attainment estimates suggest that further work is required in this area.

This paper was originally motivated by the view that weak data is likely to be one of the main reasons for the discouraging results obtained in the recent empirical literature on human capital and growth. Our results clearly support this hypothesis, as does recent work by Krueger and Lindhal (2001) and Cohen and Soto (2001), and suggest that the contribution of investment in education to productivity growth is sizable. Unlike several older data sets, our revised series produce positive and theoretically plausible results using a variety of growth specifications and, unlike MRW's original (1992) results for a similar sample, our findings survive a simple robustness check. More importantly, our analysis of the performance of different schooling data sets in a variety of production function specifications shows a clear tendency for human capital coefficients to rise and become more precise as the information content of the schooling data increases. We have extrapolated this tendency to construct estimates of the value of the coefficient that would be obtained with the correctly measured stock of human capital. Using an extension of the classical errors-in-variables model, we correct for measurement error bias and produce meta-estimates of the elasticity of output with respect to average years of schooling that suggest that the true value of this coefficient is almost certainly above 0.50.

References

- Barro, R. and J-W. Lee (1993). "International Comparisons of Educational Attainment." *Journal of Monetary Economics* 32, pp. 363-94.
- Barro, R. and J-W. Lee (1994). "Data set for a panel of 138 countries." Mimeo, Harvard University.
- Barro, R. and J-W. Lee (1996). "International Measures of Schooling Years and Schooling Quality." *American Economic Review* 86(2), *Papers and Proceedings*, pp. 218-23.
- Barro, R. and J-W. Lee (2000). "International data on educational attainment, updates and implications." NBER Working Paper no. 7911. (<http://www.nber.org/papers/w7911>).
- Behrman, J. and M. Rosenzweig (1994). "Caveat emptor: cross-country data on education and the labor force." *Journal of Development Economics* 44, pp. 147-71.
- Benhabib, J. and M. Spiegel (1994). "The Role of Human Capital in Economic Development: Evidence from Aggregate Cross-Country Data." *Journal of Monetary Economics* 34, pp. 143-73.
- Card, D. (1999). "The Causal Effect of Education on Earnings." In O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, vol. 3, chapter 30, pp. 1801-1863. North Holland.
- Caselli, F., G. Esquivel and F. Lefort (1996) "Reopening the convergence debate: a new look at cross-country growth empirics." *Journal of Economic Growth*, pp. 363-89.
- Castles, I. (1995?). *Year Book Australia 1994*. Australian Bureau of Statistics, Canberra.
- Cohen, D. and M. Soto (2001). "Growth and human capital: good data, good results." CEPR Discussion Paper no. 3025.
- Dabán, T., R. Doménech and C. Molinas (1997). "International and intertemporal comparisons of real product in OECD countries: a growth sensitivity approach," *Review of Income and Wealth* 43(1), pp. 33-48.
- de la Fuente, A. (2002a). "Human capital and growth: a survey and some tentative conclusions." Mimeo, Instituto de Análisis Económico (CSIC), Barcelona.
- de la Fuente, A. (2002b). "On the sources of convergence: A close look at the Spanish regions." *European Economic Review* 46(3), pp. 569-99.
- de la Fuente, A. and R. Doménech (2000). "Human capital in growth regressions: how much difference does data quality make?" CEPR Discussion Paper no. 2466/ Economics Department Working Paper no. 262, OECD, Paris.
- de la Fuente, A. and R. Doménech (2001). "Schooling data, technological diffusion and the neoclassical model." *American Economic Review* 91(2), *Papers and Proceedings*, pp. 323-27.
- Griliches, Z. (1980). "R&D and the productivity slowdown." *American Economic Review* 70 *Papers and Proceedings*, pp. 343-48.
- Griliches, Z. (1986). "Economic data issues." In Z. Griliches and M. Intrilligator, editors, *Handbook of Econometrics*. North Holland, Amsterdam.

- Hamilton, J. and J. Montenegro (1998). "The augmented Solow model and the productivity slowdown." *Journal of Monetary Economics* 42, pp. 495-509.
- Harmon, C., I. Walker and N. Westergaard-Nielsen (2001). "Returns to education in Europe." In *Public funding and private returns to education. PURE final report*.
- Islam, N. (1995). "Growth empirics: a panel data approach." *Quarterly Journal of Economics* 110, pp. 1127-70.
- Kaneko, M. (1986). "The educational composition of the world's population: a database." Washington DC, the World Bank, Education and Training Department, Report No. EDT 29.
- Knight, M., N. Loayza and D. Villanueva (1993). "Testing the neoclassical theory of economic growth: A panel data approach." *IMF Staff Papers* 40(3), pp. 512-41.
- Krueger, A. and M. Lindahl (2001). "Education for growth: why and for whom?" *Journal of Economic Literature* XXXIX, pp. 1101-36.
- Kyriacou, G. (1991). "Level and Growth Effects of Human Capital, A Cross-Country Study of the Convergence Hypothesis." Mimeo, NYU.
- Lau, L., D. Jamison, and F. Louat (1991). "Education and Productivity in Developing Countries: an Aggregate Production Function Approach." Report no. WPS 612, the World Bank.
- Lau, L., S. Bhalla and F. Louat (1991). "Human and physical capital stock in developing countries: construction of data and trends." Draft mimeo, *World Development Report*, World Bank.
- Lee, J-W. and R. Barro (2001). "Schooling quality in a cross-section of countries." *Economica* 68, pp. 465-88.
- Mankiw, G., D. Romer and D. Weil (1992). "A Contribution to the Empirics of Economic Growth." *Quarterly Journal of Economics*, pp. 407-37.
- Mas, M., F. Pérez, E. Uriel and L. Serrano (1998). *Capital humano, series históricas*. Forthcoming by Fundación Bancaja.
- Nehru, V., E. Swanson and A. Dubey (1995). "A New Database on Human Capital Stocks in Developing and Industrial Countries: Sources, Methodology and Results." *Journal of Development Economics*, 46, pp. 379-401.
- Nordic Statistical Secretariat (1994). *Educational indicators in the nordic countries*.
- OECD (1970a). *Statistics of the occupational and educational structure of the labour force in 53 countries*. Paris.
- OECD (1970b). *Development of higher education, 1950-67. Statistical Survey*. Paris.
- OECD (1974). *Annuaire des Statistiques de l'Enseignement, Vol. I*, Paris.
- OECD (1990). *L'Enseignement dans les Pays de l'OCDE, 1987-88. Recueil d'informations statistiques*. Paris.
- OECD (1995). *OECD Education Statistics 1985-92*. Paris.
- OECD (various years). *Education at a Glance*. Paris.
- Pritchett, L. (1999). "Where has all the education gone?" Mimeo, World Bank.
- Psacharopoulos, G. and A. Arriagada (1986). "The educational composition of the labour force: an international comparison," *International Labour Review*, 125(5), Sept.-Oct., pp. 561-74.

Steedman, H. (1996). "Measuring the quality of educational outputs: a note." Center for Economic Performance, Discussion Paper no. 302, LSE.

Summers, R. and Heston.....

Temple, J. (1998). "Robustness tests of the augmented Solow model." *Journal of Applied Econometrics* 13, pp. 361-75.

UNESCO (1993). *Statistics of Educational Attainment and Illiteracy, 1970-80*. Division of Statistics on Education. UNESCO, Paris.

UNESCO (various years). *Statistical Yearbook*, Paris.

United Nations (various years). *Demographic Yearbook*, New York.