ARTICLE IN PRESS

DEVEC-01760; No of Pages 13

Journal of Development Economics xxx (2013) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

The contribution of schooling in development accounting: Results from a nonparametric upper bound ☆

Francesco Caselli [a,*], Antonio Ciccone [b]

[a] LSE, CEP, CEPR, NBER, Centre for Macroeconomics, United Kingdom
[b] ICREA-UPF, Barcelona GSE, CEPR, Spain

ARTICLE INFO

Article history:
Received 28 November 2011
Received in revised form 2 October 2012
Accepted 9 February 2013
Available online xxxx

JEL classification:
O1
O3
O4

Keywords:
Development accounting
Imperfect substitution
Schooling

ABSTRACT

How much would output increase if underdeveloped economies were to increase their levels of schooling? We contribute to the development accounting literature by describing a nonparametric upper bound on the increase in output that can be generated by more schooling. The advantage of our approach is that the upper bound is valid for any number of schooling levels with arbitrary patterns of substitution/complementarity. Another advantage is that the upper bound is robust to certain forms of endogenous technology response to changes in schooling. We also quantify the upper bound for all economies with the necessary data, compare our results with the standard development accounting approach, and provide an update on the results using the standard approach for a large sample of countries.

1. Introduction

Low GDP per worker goes together with low schooling. For example, in the country with the lowest output per worker in 2005, half of the adult population has no schooling at all and only 5% has a college degree (Barro and Lee, 2010). In the country with output per worker at the 10th percentile, 32% of the population has no schooling and less than 1% a college degree. In the country at the 25th percentile, the population shares without schooling and with a college degree are 22% and 1% respectively. On the other hand, in the US, the share of the population without schooling is less than 0.5% and 16% have a college degree.

To some extent, such differences in attainment could reflect efficient schooling decisions in response to international differences in technology or institutional quality (e.g. Foster and Rosenzweig, 1995; Jensen, 2010; Munshi and Rosenzweig, 2006). On the other hand, it seems highly plausible that schooling attainment in poor countries is also limited by lack of access to schools (particularly in rural areas), and credit constraints that force parents to send children to work in order to provide for current consumption. Credit constraints also limit poor parents' capacity to cover tuition, uniforms, and meals. Consistent with the view that there are barriers to investment in schooling, Duflo (2001) finds large enrollment effects from an expansion in public school provision, and Schultz (2004) from the introduction of a conditional cash transfer program. The crucial importance of public funding (and other government policies) to enable mass schooling is discussed at length in Goldin and Katz (2008). It is also consistent with the existence of barriers to attainment that the returns to schooling are higher in poor countries than in rich ones (e.g. Bils and Klenow, 2000). The view that schooling attainment is in part limited by lack of access and credit constraints has led national governments, bilateral and multilateral donors, and civil-society NGOs to prioritize schooling attainment among their development goals for several decades. For example, one of the millennium development goals is universal education.

But how much of the output gap between developing and rich countries can be accounted for by differences in the quantity of schooling? Early empirical attempts to answer this question using cross-country data focused on regressions of growth (or GDP levels) on measures of educational enrollment or attainment (e.g. Barro, 1991; Benhabib and Spiegel, 1994; Caselli et al., 1996; Mankiw et al., 1992; see Krueger and Lindhal, 2001 for a survey and evaluation of this literature). One difficulty with this literature is that results on the impact of schooling did not prove robust to alternative measures of the education variable, the sample, or the estimation method. Also, it proved difficult to tackle the problem of endogeneity of schooling.

In part in response to these difficulties with the regression approach, a second wave of studies focused on calibration rather than estimation (e.g. Hall and Jones, 1999; Hendricks, 2002; Klenow

☆ We thank Marcelo Soto, Hyun Son, the two anonymous referees, and very especially David Weil for their useful comments.

* Corresponding author. Tel.: +44 207 955 7498.
E-mail addresses: f.caselli@lse.ac.uk (F. Caselli), antonio.ciccone@upf.edu (A. Ciccone).

0304-3878/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jdeveco.2013.02.006

and Rodriguez-Claire, 1997), giving rise to a thriving new literature known as development accounting. A robust result in the development accounting literature is that only a relatively small fraction of the output gap between developing and rich countries can be attributed to differences in the quantity of schooling.[1] This result appears to dampen expectations that current efforts at boosting schooling in poor countries, even if successful, will do much to close the gaps in living standards.[2]

The somewhat negative result from development accounting is obtained using a parametric approach. Technology differences across countries are assumed to be skill neutral, and workers with different attainment are perfect substitutes. Relative wages are then used to gauge the relative efficiency in production of workers with different attainment. A potential concern is that there is by now a consensus that differences in technology across countries or over time are generally not Hicks-neutral, and that perfect substitutability among different schooling levels is rejected by the empirical evidence (e.g. Angrist, 1995; Autor and Katz, 1999; Caselli and Coleman, 2006; Ciccone and Peri, 2005; Goldin and Katz, 1998; Katz and Murphy, 1992; and Krusell et al., 2000). Once the assumptions of perfect substitutability among schooling levels and Hicks-neutral technology differences are discarded, can we still say something about the output gap between developing and rich countries attributable to schooling?

Answering this question while sticking to a parametric approach requires assuming that there are only two imperfectly substitutable skill types, that the elasticity of substitution between these skill types is the same in all countries, and that this elasticity of substitution is equal to the elasticity of substitution in countries where instrumental-variable estimates are available (e.g. Angrist, 1995; Ciccone and Peri, 2005). These assumptions are quite strong. For example, the evidence indicates that dividing the labor force in just two skill groups misses out on important margins of substitution (Autor et al., 2006; Goos and Manning, 2007). Once there are more than 3 skill types, estimation of elasticities of substitution becomes notoriously difficult for two main reasons. First, there are multiple, non-nested ways of capturing patterns of substitutability/complementarity and this make it difficult to avoid misspecification (e.g. Duffy et al., 2004). Second, relative skill supplies and relative wages are jointly determined in equilibrium and estimation therefore requires instruments for relative supplies. It is already challenging to find convincing instruments for two skill types and we are not aware of instrumental-variable estimates when there are 3 or more imperfectly substitutable skills groups.

We explore an alternative to the parametric production function approach. In particular, we make the observation that when aggregate production functions are weakly concave in inputs, assuming perfect substitutability among different schooling levels yields an upper bound on the increase in output that can be generated by more schooling. This is true irrespective of the pattern of substitutability/complementarity among schooling levels, as well as the pattern of cross-country non-neutrality in technology. This basic observation does not appear to have been made in the development accounting literature. It is worthwhile noting that the production functions used in the development accounting literature satisfy the assumption of weak concavity in inputs. Hence, our approach yields an upper bound on the increase one would obtain using the production functions in the literature. Moreover, the assumption of weakly concave aggregate production functions is fundamental for the development accounting approach as it is clear that without it, inferring marginal productivities from market prices cannot yield interesting insights into the factors accounting for differences in economic development.

The intuition for why the assumption of perfect substitutability yields an upper bound on the increase in output generated by more schooling is easiest to explain in a model with two schooling levels, schooled and unschooled. In this case, an increase in the share of schooled workers has, in general, two types of effects on output. The first effect is that more schooling increases the share of more productive workers, which increases output. The second effect is that more schooling raises the marginal productivity of unschooled workers and lowers the marginal productivity of schooled workers. When assuming perfect substitutability between schooling levels, one rules out the second effect. This implies an overstatement of the output increase when the production function is weakly concave, because the increase in the marginal productivity of unschooled workers is more than offset by the decrease in the marginal productivity of schooled workers. The result that increases in marginal productivities produced by more schooling are more than offset by decreases in marginal productivities continues to hold for an arbitrary number of schooling types with any pattern of substitutability/complementarity as long as the production function is weakly concave. Hence, assuming perfect substitutability among different schooling levels yields an upper bound on the increase in output generated by more schooling.

From the basic observation that assuming perfect substitutability among schooling levels yields an upper bound on output increases, and with a few ancillary assumptions – mainly that physical capital adjusts to the change in schooling so as to keep the marginal product of physical capital unchanged – we derive a formula that computes the upper bound using exclusively data on the structure of relative wages of workers with different schooling levels. We apply our upper-bound calculations to two data sets. In one data set of 9 countries we have detailed wage data for up to 10 schooling-attainment groups for various years between 1960 and 2005. In another data set of about 90 countries we use evidence on Mincerian returns to proxy for the structure of relative wages among 7 attainment groups. Our calculations yield output gains from reaching a distribution of schooling attainment similar to the US that are sizeable as a proportion of initial output. However, these gains are much smaller when measured as a proportion of the existing output gap with the US. These results are in line with the conclusions from development accounting (e.g. Caselli, 2005; Hall and Jones, 1999; Klenow and Rodriguez-Claire, 1997). This is not surprising as these studies assume that workers with different schooling attainment are perfect substitutes and therefore end up working with a formula that is very similar to our upper bound.

A potential limitation of the parametric approach to development accounting is that it typically assumes that changes in schooling attainment leave technology unchanged.[3] This assumption would be wrong if there were important schooling externalities or significant appropriate-technology effects. We discuss the extent to which our nonparametric upper bound is robust to endogenous technology

[1] Recently this result has been challenged by Gennaioli et al. (forthcoming), who argue that much of top managers' and entrepreneurs' returns to schooling are formally earned as profits, and therefore unaccounted for by standard microeconomic estimates of the returns to schooling – a key ingredient in most development-accounting calculations. After accounting for managers' returns to schooling, they argue that the average Mincerian return to schooling is around 20%, about double what is usually found in the literature. Using this higher return leads to a large increase in the explanatory power of human capital for income differences. Gennaioli et al.'s estimate of managers' returns to schooling is based on firm-level valued-added regressions that do not control for manager characteristics other than schooling. As such characteristics may be correlated with managers' schooling, it is difficult to know what part of the return can be attributed to schooling only.

[2] Partially in response to these findings, some authors have advocated a shift to cross-country differences in the quality of schooling (e.g. Erosa et al., 2010; Hanushek and Woessmann, 2008, 2011; Manuelli and Sheshadri, 2010). Other authors have emphasized aspects of human capital such as health (Weil, 2007) and experience (Lagakos et al., 2012).

[3] This is not always the case however. For example, a recent paper by Jones (2011) computes rich–poor human capital ratios using relative wages in poor as well as rich countries. His approach implies that computed human capital ratios will also reflect differences in human capital quality and – to the extent they affect relative wages – differences in technology. In Jones' framework, the perfect substitution case yields a lower bound on the income increase that can be achieved by raising human capital in poor countries.

Please cite this article as: Caselli, F., Ciccone, A., The contribution of schooling in development accounting: Results from a nonparametric upper bound, Journal of Development Economics (2013), http://dx.doi.org/10.1016/j.jdeveco.2013.02.006