

Education and economic growth

Jonathan Temple*

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN

1. Introduction

This note provides a brief overview of the effects of education on productivity, drawing on both theoretical work and empirical evidence. The main focus is on the available evidence, and in particular what it tells us about the private and social returns to education. The focus on these returns is motivated by an observation central to economics, namely that it is where private and social returns diverge that there may be a case for using government taxes and subsidies, or other policy instruments, to realign incentives.

As we will see in what follows, there are several instances where an individual undertaking education may not take into account all the costs and benefits of that education to society as a whole. These 'externalities' are a form of market failure. They offer one of the best reasons for governments to be wary of approaches to education that are based entirely on decentralised, market-based provision.

Given the limited space available, the paper is at best a selective overview. For example, the paper does not cover the effects of training and vocational education. The forms of evidence I know best, and draw on here, are much better equipped to measure the productivity effects of education than those of training. This is a potentially important omission, because historical accounts of Britain's productivity gap have often emphasised a perceived weakness in training provision, with consequences for the distribution of income as well as productivity.

Another omission is any consideration of the wider benefits of education, beyond productivity. Again this is a choice motivated partly by the available evidence, rather than a judgement that these benefits are unimportant. Nevertheless, some may find this omission puzzling. After all, the traditional case for education is cast in terms of personal development, to the advantage of individuals and their society as a whole. Why this paper's narrow focus on productivity?

One answer might run as follows. When economists analyse the effects of education on labour productivity, they are sometimes drawn to the conclusion that these narrow effects are large enough by themselves to justify current or increased expenditure, with some margin for error. On this view, learning about productivity effects can inform decision-making, whatever the perspective one takes on the wider value of education.

In discussing further the effects of education on growth, I will generally continue to sound this note of cautious optimism. Yet there is at least one important qualification, and one that might be regarded as uncontroversial. *Whatever education does for growth, it is unlikely to do quickly.* The reason is that changes to educational provision will typically affect only new entrants to the labour force, and these new entrants are only a small fraction of all those currently employed. This suggests that the consequences of a shift in education policy will emerge fully only over several decades; education is rarely, perhaps never, a quick fix.

The remainder of the paper is structured as follows. In the next section, I briefly discuss the perspective of recent work on growth theory, and the implications for education policy. Section 3 reviews the key part of the existing statistical evidence, based on survey data on wages and education. Section 4 will cover the evidence from comparing growth performance across countries. Section 5 briefly discusses the implications of growth theory for education spending, before section 6 concludes.

2. The perspective of growth theory

The models of new growth theory, or 'endogenous growth', see productivity improvements as driven by the creation and implementation of new ideas. One theoretical approach is to analyse the incentives for R&D, so that the rate of technological progress is determined within the model, hence the term 'endogenous'. Since economists might be able to say something useful about how R&D is determined, these models potentially allow some insight into the determinants of long-run growth rates.

What are the implications for education policy? Theoretical models usually assume that R&D activities, and thus long-run growth, depend on the availability of skilled labour. The models allow one to analyse the extent to which skilled individuals will conduct research, given the presence of other job opportunities. The usual conclusion is that, when a relatively large number of skilled people are available, more become engaged in R&D activities, more ideas are generated or implemented, and this raises the long-run growth rate.¹

The argument does not rest on any assumptions about the role of universities as themselves providers of research. In the framework of new growth theory, education is important because some of the educated will create new ideas in the private sector, not because the educated engage in university research. Research funded by the state may be an important complement to privately funded R&D, but the claim that education matters for growth has greater generality.

It is also clear that individuals are unlikely to capture all the benefits of their research; for example, the knowledge generated by one firm's research department may help another firm working on a different product. These knowledge spillovers are a form of externality. From society's point of view, individuals are likely to under-invest in certain kinds of education, because they do not value the wider benefits of their later research. This argument starts to provide a case for a selective subsidy, targeted at subjects that can provide a foundation for later work in research and development. To put this in more concrete terms, if one were to rely solely on the market, there might be a tendency to train too many lawyers and too few engineers.

One of the leading figures in the new growth theory, Paul Romer, has recently discussed how these models might inform education policy. He notes that, in the models reviewed above, growth is

* This is a paper prepared following my presentation at the HM Treasury seminar on economic growth and government policy, 12 October 2000. It provides a brief summary of the existing theory and evidence on the role of education in productivity growth.

¹ The most obvious mechanism is that, with a relatively large pool of skilled people to draw on, the costs of employing researchers and carrying out R&D will be lower than otherwise. But the mechanisms and implications of these models are richer than this necessarily short summary implies.

determined by the quantity of inputs used in R&D, not simply expenditure upon it. One reason this point matters is that incentives to encourage R&D, such as tax credits, may be ineffective unless they encourage a greater number of scientists and engineers to work towards developing new ideas.

To illustrate this, consider a very simple model, in which a fixed supply of scientists only work in R&D and are the only input to the research process. Then an increase in R&D spending will simply raise the wages of scientists, with no effect on the number of researchers engaged in R&D or the growth rate.

In a more general and realistic model, there will be some effect of greater R&D spending on total research inputs and therefore growth.² To create a large effect, higher wages for scientists should encourage more individuals to train as scientists. This requires some flexibility on the part of the education system, and the provision of relevant information to potential students. So the effectiveness of direct subsidies or tax credits for R&D may be enhanced by complementary education policies, aimed at improving or subsidizing the supply of research inputs, rather than simply the demand for them.

Some related work, admittedly more speculative, also emphasises the possible interdependence of decision-making by firms on R&D, and decisions by individuals to acquire knowledge and skills. One possible result is a 'low-skill, low-quality trap', whereby low skill levels and slow rates of innovation could reflect a coordination failure. Individuals don't invest in education, because they do not expect firms to make use of the resulting skills; while firms don't innovate rapidly, because they do not expect to have the requisite skills base to draw upon. It is easy to see how this situation could become self-reinforcing, as it does in some of the models; hence the term 'trap'.

Taken together, the models of the new growth theory are important for several reasons. They see knowledge and skills as a key input in the creation of new ideas. This mechanism provides the most plausible justification for viewing education as a central determinant of growth rates over long time intervals. Some theoretical work also points to specific ways in which laissez-faire, market-based outcomes might be improved upon. Overall, the models suggest that, in searching for ways to raise productivity growth, policy on education might be one of the first places to look.

3. Evidence from labour economics

Most of the remainder of the paper turns to evidence on education and growth. The most convincing empirical work is based on surveys of thousands of individuals, and uses this survey data to estimate statistical models for earnings. These models typically explain earnings differentials across individuals using educational attainment, a measure of labour market experience, and other observable characteristics such as gender. The idea is to compare earnings differentials for people who differ by educational attainment, but are alike in other respects. Since an individual's earnings will often be related to their productivity, this correlation may tell us something useful about the productivity effects of education, and is a natural place to start.³

Indeed it is possible to go further: one can treat an individual's educational choices as if these represent investment decisions and, under some assumptions, derive a measure of the private rate of return to an extra year of schooling. Most calculations assume that the main cost to the student is forgone earnings, and the main benefit an increase in earnings later in life, although the form of approach is more general. Studies of this kind almost always

indicate that the private rate of return to an extra year of schooling is in the region of 5% to 15%. This probably compares quite favourably with the other investment opportunities open to most individuals.

The initial evidence, of the 1960s and 1970s, was criticised on the grounds that the education-earnings correlation might be spurious. The reason is that high levels of educational attainment will be associated with unobservable characteristics (ability, ambition, determination and so on) that also have implications for earnings. More recent work has found a number of ingenious ways around this problem, primarily by seeking out 'natural experiments' that allow more confident conclusions. For instance, one can look at the earnings differentials of twins who differ in educational attainment. Another good example of the 'natural experiment' approach is a paper using the fact that, because compulsory schooling laws sometimes allow pupils to leave school on reaching a certain age, the duration of schooling will have some tendency to vary with an individual's quarter of birth. This can be used to isolate variation in educational attainment that is plausibly independent of other characteristics; it is therefore less vulnerable to the charge that the correlation between education and earnings simply reflects omitted variables like ability.

Yet there remain some widely acknowledged limitations to this form of approach. The empirical work probably measures private returns quite accurately, but for most questions of interest to policy-makers, it is the social return that is critical.⁴ A particular concern is that educational qualifications may be valued in the labour market solely because they act as a signal of ability (or other characteristics, like ambition) and not because more educated individuals are genuinely more productive. This means that the social return could be less than the private return. There is almost certainly some truth to the signalling view, but precise evaluation remains difficult.

There are also one or two good reasons to believe that social returns are above, not below, private returns. The potential importance of externalities has already been discussed. A more specific idea is that education may play a role as a filtering mechanism, which allows a more efficient matching of workers to jobs than would otherwise be possible. This improvement in matching could raise overall productivity, without these gains being captured by the workers themselves.

4. Evidence across countries

The potential importance of signalling and filtering effects, and externalities, should be clear. One relatively simple way to assess their importance might be to compare the economic performance of countries with different levels of educational provision, while controlling for other factors. Another advantage of this approach, compared to that based on microeconomic data, is that it may be better equipped to assess the predictions of growth models.

Before going down this route, we might want some evidence on whether education could ever be important enough to have an impact discernible at the aggregate level. Using simple models, we can examine the contribution of international variation in educational attainment to international differences in living standards. One recent exercise of this kind takes the estimates of private returns described above at face value, and equates them with social returns. Under this admittedly very strong assumption, education can explain nearly all the variation in living standards across OECD members.

As with the theoretical work reviewed in section 2, this indicates that education could easily matter enough for us to detect effects

2 For example, some individuals outside the R&D sector, but originally trained as scientists, may switch into R&D careers in response to higher wages.

3 Although it will not necessarily help us evaluate the arguments of the models described in section 2. That may require cross-country evidence, of the kind reviewed in section 4 below.

4 Occasionally papers present figures for something they call the social return, but in practice this is usually just the private return adjusted for the use of resources by the education sector and tax payments. This may be closer to the relevant quantity, but remains only a proxy for it.

in statistical models, based on comparing growth performance across countries. Such exercises no longer take the private returns at face value, but instead provide (or aim to provide) an independent measure of the importance of education. How well do the estimates of educational effects at the macro level tally with the micro evidence?

At least until recently, the answer has been – not very well. There are several studies that report a substantial impact of education on growth, but these are based on large samples of countries mostly drawn from the developing world. Relatively few papers offer much insight into the questions relevant to the UK. There are even some well-known studies that use the macroeconomic evidence to question whether education matters at all.

This evidence has to be interpreted carefully, since modelling the variation of growth across countries is not an easy task. The researcher's aim to control for other relevant factors has to contend with data of somewhat dubious quality, with considerable uncertainty about the appropriate form of the statistical model, and samples of countries that are often too few in number to yield unambiguous results. For these reasons, the area is something of a minefield, and it would be fair to say that most economists place rather more faith in the microeconomic evidence reviewed earlier.

Yet, perhaps more promisingly, recent work based on cross-country evidence has started to yield results that look more consistent with the micro data. One study restricted to the OECD countries takes care to adjust the education data for various measurement problems, and finds stronger results. Another approach is to implement statistical methods that adjust for measurement error. Again, this tends to strengthen the measured impact of education.

The field is still developing, and even the most recent results should be treated with caution. It is quite conceivable that the associated statistical problems will never be resolved to the satisfaction of reasonable observers. My own, perhaps idiosyncratic, reading of the cross-country literature is that it is starting to provide tentative support for the traditional interpretation of the microeconomic evidence. To put this another way, it has become harder to argue that the private rates of return estimated from micro data reflect only signalling effects; but perhaps only slightly.

5. Unbalanced growth and why it matters

This section departs slightly from the main themes of the paper, to examine the implications for education spending of models of 'unbalanced growth'. The aim is to convey a simple but quite powerful argument that is widely accepted by economists, but probably not well understood by some commentators.

The starting point is to think of the economy in terms of two sectors: one that provides education, and another providing all other goods and services. The main costs of the education sector are the salaries of teachers. The final key assumption is that, at least over long spans of time, one might expect productivity growth to be lower in the provision of education than in the rest of the economy.

Does this matter? The answer is yes, a lot. Productivity growth in the rest of the economy will inevitably drive up real wages. This will make teaching less attractive as a career, and ultimately the only way to overcome this is to raise real wages for teachers, in line with the increases elsewhere in the economy. The immediate implication is that education spending must grow in real terms, year on year, simply to maintain a constant level and quality of provision.⁵

The alternative is to allow the relative wages of teachers to fall, in which case one will tend to see a gradual decline in the talents of

those entering the profession, given that people can do other things. There are some observers who think this process is happening in Britain's universities, where departments in certain subjects are finding it increasingly difficult to retain those whose skills and talents are highly valued in the labour market, as in subjects like law and engineering.⁶

In the short run, this is not necessarily a concern. The optimal degree of talent for teachers and lecturers is not an easy question to settle, and there should not be a presumption that the optimal degree of talent is greater than the level presently found. Yet it should also be clear that the strategy of allowing relative wages to decline is not sustainable indefinitely: if unchecked, it would eventually lead to a position where sixth-formers are taught by school-leavers.

These conclusions are less alarming than they might appear. Recall that the argument is predicated on productivity growth elsewhere in the economy. The same productivity growth that is driving up the real wages of teachers also makes available the extra resources needed to raise educational spending, with room to spare. Nevertheless, the argument is a useful reminder that a constant level of education spending, adjusted for general inflation, may buy less and less education over time – however paradoxical that may seem at first.

6. Conclusions

In the limited space available here, it is possible only to hint at the full extent of available research on education and growth. I have tried to convey the main findings, at considerable risk of oversimplification.

First of all, there are well-established theoretical reasons to believe that education plays a major role in productivity improvement. The models that yield this proposition are based on mechanisms that are relatively plausible and command wide agreement among economists.

Secondly, and perhaps somewhat more controversially, there are also some reasons to believe that laissez-faire or market-based outcomes can sometimes be improved upon. In thinking about this claim further, the distinction between private and social returns to education is central. Most studies imply that the private returns to education are high, but even the most sophisticated papers are vulnerable to the argument that such returns depend on signalling effects. Working in the other direction, the presence of externalities and filtering suggests that the estimated private return could understate the true social return. Empirical work comparing experiences across countries may eventually offer some insight into these questions, but reliable answers are some way off.

Theoretical work also has a contribution to make, and even quite simple ideas and analysis can offer some useful insights. Potential examples include the above discussions of the implications of slow productivity growth in the education sector, and the case for education policies that complement direct subsidies or tax credits for R&D.

More recent research on education by economists is beginning to examine other issues, notably the institutional structure through which education is provided. It is possible that more work of this kind, both theoretical and empirical, could yield substantial benefits for policy.

Overall, the case for productivity effects of education is not yet proven, at least not beyond reasonable doubt. Nevertheless, one could also find some casual evidence in favour of the importance of education. In searching for at least a partial explanation for the recent dynamism of the US, and its leading position in new

5 In talking about growth 'in real terms', here and later on, I mean that nominal expenditure on education must rise more quickly than a price deflator for the whole economy. One could recast the argument made here by saying that a price deflator constructed specifically for the costs of education will tend to show more rapid inflation than found in the economy as a whole. Also, note that I am holding constant other factors, including demographic changes.

6 Some of the relevant issues and evidence are discussed in more depth in Greenaway and Haynes (2000).

technologies, one candidate would be the strengths of its universities. In this respect it is interesting to note that in 1995 the US spent 2% of its GDP on tertiary education (including both private and public spending, and university research). This compares to 1.1% of GDP in the EU, and 0.7% of GDP in the UK.

Appendix and references

This short overview draws on a much longer and more technical review of the literature, Temple (2001), prepared for the OECD as part of their Growth Project. An excellent review of microeconomic estimates of the returns to schooling, with a focus on the UK, can be found in Harmon et al. (2000).

The implications of R&D based growth models for education policy, including specific policy proposals for the US, are discussed in Romer (2000).

The paper alluded to at the start of section 4 above is Woessmann (2000).

The figures on spending in the conclusions (section 6) are based upon OECD sources, as cited in the Greenaway-Haynes report (p. 26-28) commissioned by the Russell Group of universities, on the future funding position of universities in the UK. Note that international differences in the research roles of the tertiary sector will undermine the usefulness of these figures for comparing teaching provision.

Greenaway, D. and Haynes, M. (2000). Funding universities to meet national and international challenges. School of Economics Policy Report, Nottingham University.

Harmon, C., Oosterbeek, H. and Walker, I. (2000). The returns to education: a review of evidence, issues and deficiencies in the literature. Centre for the Economics of Education, London School of Economics, December.

Romer, P. M. (2000). Should the government subsidize supply or demand in the market for scientists and engineers? *NBER working paper no. 7723*.

Temple, J. R. W. (2001). Growth effects of education and social capital in the OECD countries. Available at <http://www.nuff.ox.ac.uk/Users/Temple/research.htm>

Woessmann, L. (2000). Specifying human capital: a review, some extensions, and development effects. The Kiel Institute of World Economics working paper no. 1007.