Volume Title: Growth and Productivity in East Asia, NBER-East Asia Seminar on Economics, Volume 13

Volume Author/Editor: Takatoshi Ito and Andrew K. Rose, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-38680-5

Volume URL: http://www.nber.org/books/ito_04-2

Conference Date: June 20-22, 2002

Publication Date: June 2004

Title: Ideas and Education: Level or Growth Effects and Their Implications for Australia

Author: Steve Dowrick

URL: http://www.nber.org/chapters/c10743

1

Ideas and Education

Level or Growth Effects and Their Implications for Australia

Steve Dowrick

Introduction

The importance of human capital for economic growth was highlighted in much of the "new growth theory" that came to prominence in the late 1980s and early 1990s. The neoclassical growth model, formalized three decades earlier, had focused on the accumulation of machinery and equipment and emphasized the feature of diminishing returns—which implied that such investment would not be able to drive long-run growth. The new generation of studies switched attention to the accumulation of human capital and the possibility that returns to investment in education, training, and research may not suffer from diminishing returns.

There is an important distinction between embodied and disembodied human capital. Human capital in the form of abilities and skills is embodied inasmuch as it lives and dies with particular people. We invest in human capital not only through formal education and training programs but also through experience on the job and through domestic and social interaction. The time and effort devoted to parenting, for example, represents an enormous investment in the human capital of the next generation.

The accumulation of abilities contributes both to psychic rewards and to marketed economic activity. Whereas the value of the former is hard to measure, there are relatively straightforward ways for us to measure the latter. Economists are only just beginning to address seriously the task of

Steve Dowrick is professor of economics and ARC Senior Research Fellow in the School of Economics, Australian National University.

The author thanks Creina Day and Shebnem Pollack for research assistance and acknowledges the Melbourne Institute for assistance in preparing the background papers for their 2002 Economic and Social Outlook Conference, and acknowledges the Australian Research Council for financial assistance.

evaluating nonmarket activities such as domestic labor—see, for example, Folbre and Nelson (2000) and Apps and Rees (2001). I focus in this paper on market-related returns to human capital because we do not yet have internationally standardized valuations of nonmarket activities.

The most extensively documented feature of embodied human capital is the relationship between education and wages. Studies of earnings in advanced capitalist economies typically find that each extra year of schooling raises earnings by 5 to 10 percent. These findings are confirmed by Australian studies such as those by Miller, Mulvey, and Martin (1995), who analyze earnings of twins and find that the return to a year of education lies between 4.5 percent and 8.3 percent, and Preston (1997) who reports high rates of return to advanced educational qualifications. The results of Miller, Mulvey, and Martin are particularly interesting because they control for the influences of genetic and domestic background to identify the direct contribution of education—following studies by Ashenfelter and Krueger (1994) and Ashenfelter and Rouse (1997, 1998) that estimate U.S. rates of return between 9 and 16 percent.

This evidence leads us to expect that, if the average educational attainment of the working-age Australian population were to rise by one year, real gross domestic product (GDP) should rise by up to 8 percent. This increase in the level of GDP will, typically, take place gradually. An increase in the length of schooling of teenagers will only increase the average educational experience of the adult population as the new, better-educated cohorts enter the workforce, replacing older cohorts. We expect the transition to last four decades, if people enter the labor force aged twenty and exit at an age of about sixty. If this is so, the annual growth rate of GDP will be 0.2 percentage points above trend during the transition period, resulting in an overall 8 percent increase, after which time the growth rate will revert to trend—with, perhaps, some lagged adjustment to the stock of physical capital. In this sense, changes in educational investment are predicted to have *growth effects* in the short run (albeit a short run of forty years), but only *level effects* in the long run.

This is the conventional approach, which treats human capital as an investment good in much the same way as a farmer might consider investing in tractors. There are, however, features of human capital that can give it a much more important role in economic development. This is particularly true when we turn our attention to disembodied human capital, the realm of knowledge and ideas that do not live and die with their inventors but can be transmitted freely between people and carried forward over generations.

A crucial economic attribute of disembodied human capital, highlighted in recent models of endogenous growth, is that ideas are both nonrival and cumulative. Nonrivalry implies that once the idea of using electronic circuits to carry out binary computations has been announced, people can simultaneously use this idea to develop a wide range of applications. One person's use of the idea does not prevent another person from using it at the same time. Moreover, ideas are cumulative: The idea of electronic computing has lead to the idea of quantum computing, which may in turn lead to yet further ideas.

Analysis of these attributes of nonrivalry and cumulative feedback has led growth theorists to speculate that investment in the generation of ideas can be the engine of long-run growth. The nonrivalry of knowledge also leads us to expect market failure. When others reap the benefits of someone's new ideas, market forces alone are unlikely to generate the optimal level of investment in knowledge—implying a need for government subsidy.

If the generation of disembodied human capital—ideas/technology—is the engine of growth, we should expect to find that embodied human capital—skills and abilities—also affect long-run growth. Ideas do not reproduce themselves without the input of highly skilled researchers. Perhaps of equal importance, the more skilled the workforce, the better it is able to absorb, implement, and adapt the new ideas emanating from the research and development (R&D) sector. To the extent that technological change is endogenous, we expect educational attainment to have long-run growth effects in addition to the conventional prediction of level effects.

In the following sections I review and evaluate evidence from recent theoretical and econometric studies relating economic growth to investment in both embodied and disembodied human capital. I restrict my attention on the empirical front to the relatively well-documented areas of investment in formal schooling and R&D, noting that this omits potentially important areas of investment in health and in informal education and training that takes place within the family and within the workplace.

1.1 Rethinking Economic Growth: The Role of Knowledge

Knowledge is fundamental to economic progress. Our material standard of living would be reduced to unrecognizable levels if we were to suffer collective amnesia—forgetting that a circular shape reduces friction, not remembering how to read and write, losing all knowledge of electrodynamics. All economic activities depend on institutions that encourage the preservation, transmission, and development of knowledge.

This seems blindingly obvious. Yet for several recent decades, the economic analysis of growth was dominated by an approach that sidelined the role of knowledge. Economists concentrated on the accumulation of objects rather than the accumulation of ideas.

The object-oriented approach to economic growth was formalized in 1956 by two economists operating at opposite ends of the globe: Robert Solow at MIT in Cambridge and Trevor Swan at ANU in Canberra. Their neoclassical growth models were formulated independently but in broadly the same way, leading to similar conclusions. Accumulation of capital—machinery, buildings, equipment, and the like—is the engine of growth in the short run. Policies that increase the share of resources going to investment will raise the productive capacity of the economy. But as the growth of the capital stock outpaces the limited resources of land and labor, the impact of each successive unit of investment is diminished. However large the boost to the investment rate, growth will eventually revert to some fixed rate determined by exogenous technological progress.

This implication of the neoclassical growth model is illustrated in figure 1.1. A boost to investment at time T_0 raises the rate of growth (the slope of the logarithmic output line) from the solid line A to the dashed line B. Ultimately, however, growth reverts to the exogenous rate, where line B becomes parallel to line A, albeit with output and incomes at a higher level than would have obtained at the lower investment rate. Tax incentives, or other policies that influence investment, affect only the level of output, not the long-run rate of growth.

The key to this conclusion is the assumption of diminishing returns to capital accumulation. Underlying this notion is the idea of capital as a collection of similar objects. A self-employed dressmaker who purchases his first sewing machine will register a large increase in annual output. Purchase of a second machine will reduce the amount of downtime when the first machine is under repair, but the consequent addition to annual output is relatively small. A third machine would probably be redundant. This assumption about diminishing returns is typically captured in growth models by postulating an aggregate production function of Cobb-Douglas

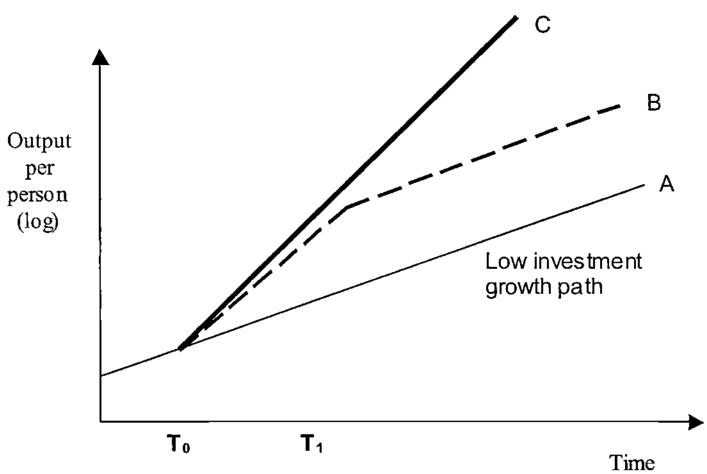

Fig. 1.1 The impact of increased investment in the neoclassical and endogenous growth models

form exhibiting constant returns to scale, where output per unit of labor at time t, y_t, is related to the net capital stock per unit of labor, k_t, as:

$$y_t = A_t(k_t)^\alpha \quad (1)$$

The elasticity of output with respect to capital, represented by the parameter α, is assumed to be less than unity. The parameter A_t represents the level of technology at time t, sometimes referred to as total factor productivity.

The marginal product of capital is

$$\frac{\partial y_t}{\partial k_t} = \frac{\alpha A_t}{k_t^{1-\alpha}}, \quad (2)$$

which, given $\alpha < 1$, diminishes toward zero as capital intensity increases.

1.1.1 The Revolution in Growth Theory: Endogenous Growth

This way of thinking about economic growth was challenged in a series of papers, starting with Paul Romer in 1986, heralded as "the new growth theory" or "endogenous growth theory."[1] A prominent feature of this new wave of economic models—indeed, their defining feature—is that policy intervention and the nature of institutions can influence the long-run growth rate of the economy.

In terms of figure 1.1, the new models suggest that policy or institutional change, instituted at time T_0, could permanently alter the slope of the growth path, as illustrated by the dotted path C.

There are various technical features of these models that make it feasible for the long-run growth rate to be determined endogenously—that is, determined by economic behavior that is analyzed within the model. One possibility arises where the degree of substitutability between capital and labor is sufficiently high that returns to the accumulation of capital do not diminish to zero.[2] We can imagine that this might be the case in some manufacturing processes where human labor is readily replaced by robots, or in the delivery of some financial services such as ATM banking. But it is not clear that this robotic model of growth is applicable to all sectors of the economy.

More interesting, to my mind at least, are models of endogenous growth that build on the economic properties of complementarity, dynamic feedback, and nonrivalry in investment. These are the properties that distin-

1. The key papers are Paul M. Romer (1986, 1990), Lucas (1988), Rebelo (1991), and Aghion and Howitt (1992). Paul Romer (1993) acknowledges the intellectual debt due to Adam Smith, Joseph Schumpeter, Arthur Lewis, and others. Further important contributions, analyzing specialization, have come from Australian economists: Yang and Borland (1991), Borland and Yang (1992), and Shi and Yang (1995).

2. This possibility was canvassed by the Australian economist John Pitchford (1960), who illustrated his argument using a constant elasticity of substitution production function.

guish the accumulation of ideas and skills from the accumulation of objects. It is worthwhile considering each of them in turn.

1.1.2 Complementarity of Investment

Complementarity arises when your investment increases the return (monetary and/or psychic) to my investment. This may occur when we invest in activities that exhibit network externalities. Learning to play chess, to speak Esperanto, or to read and write becomes much more rewarding for me if others invest in the same skills. Complementarity is not exclusive to investment in human capital; the benefits I get from investing in a telephone line and a fax machine are also enhanced when others do the same. But complementarity is probably more pervasive in the accumulation of skills than in the accumulation of objects. Indeed, such complementarity is an essential ingredient of the development of "social capital."

Complementarity is a feature of the endogenous growth model of Lucas (1988), where the productivity of any worker is enhanced not only by his or her individual level of skill but also by the average skill level among their fellow workers. This implies that the economic analysis of external effects is relevant to growth. Although my productivity depends in part on your human capital, I cannot expect you to take that into account when you decide how much education and training to undertake—and vice versa. So if we make individual decisions about the time and money we spend on education and training, we are likely to underinvest. It follows, from Lucas's analysis of such externalities, that there may be an important role for government to play. Subsidizing education will improve economic welfare in the sense that everyone will be better off as a result of an increase in human capital.[3]

1.1.3 Dynamic Feedback

These education externalities are not, however, sufficient in themselves to drive long-run growth. In Lucas's model, the rate of output growth is still limited by diminishing returns to the accumulation of both physical and human capital. He endogenizes growth by appealing to another feature of education: dynamic feedback. As we learn more, it becomes easier to acquire further knowledge and skills. An obvious example is reading. Once we have learned this skill, the acquisition of further information and skills is facilitated through book learning.

This view of dynamic feedback can be represented by a function expressing the change in the level of human capital in some representative

3. This is not the only reason for subsidizing education. Given that many parents are constrained in financing their children's education, there are both equity and efficiency reasons for public support.

household as a function of the amount of adult labor time, L_h, that is devoted to education (of self or of children) and the current level of human capital per person, h_t.

$$\frac{dh_t}{dt} = \phi \cdot L_h \cdot h_t^{\gamma} \tag{3}$$

The extent of dynamic feedback is captured by the value of the exponential parameter γ. A value of zero implies that there is no feedback. Aggregate output per person, y, now depends on both physical and human capital per person:

$$y_t = A(k_t)^{\alpha}(h_t)^{\beta}, \tag{4}$$

where we maintain the assumption of diminishing returns by restricting α, $\beta < 1$.

Endogenous growth is made feasible by the existence of positive feedback in the second sector of this economy, the education sector. To demonstrate this, take logarithms of equation (4), differentiate with respect to time and substitute equation (3) to derive the growth rate of output per worker:

$$\frac{dy_t}{dt} \cdot \frac{1}{y_t} = \alpha \frac{dk_t}{dt} \cdot \frac{1}{k_t} + \beta \frac{dh_t}{dt} \cdot \frac{1}{h_t} = \alpha \frac{dk_t}{dt} \cdot \frac{1}{k_t} + \beta \frac{\phi L_h}{h_t^{1-\gamma}} \tag{5}$$

Whether or not the accumulation of human capital can drive long-run growth is determined by the final term in this equation. With no positive feedback (i.e., if $\gamma = 0$), this term diminishes to zero as the level of human capital, h_t, increases over time. (This is exactly what happens to the physical capital term, as a given investment rate leads to slower and slower proportional growth in the stock.) But if there is sufficiently high feedback in human capital accumulation (i.e., if $\gamma = 1$), the final term in equation (5) is a positive constant. That is to say, the long-run growth rate is positive. Moreover, it is increasing in the amount of labor time that is devoted to education.

Given sufficient dynamic feedback, public subsidy of education and training can increase long-run growth. In the presence of positive externalities, or other sources of market failure, such policy will also increase economic welfare.

1.1.4 Embodied or Disembodied Human Capital

Is it reasonable, however, to suppose that the feedback effect is sufficiently strong to make education the engine of long-run growth? Note that even if the feedback parameter is close to unity—say, $\gamma = 0.9$—the long-run rate of growth in equation (5) will diminish to zero as the level of human capital increases. Stable long-run growth requires a parameter value of unity. It also requires that there be no limit to the accumulation of human capital. Human capacities to think, organize, and remember are, however, usually presumed to be finite. Moreover, our skills and abilities die with us and have to be replaced in every successive generation.[4]

In addressing the problem of limits to human capabilities, Paul M. Romer (1990) emphasizes the distinction between the skills and abilities that are embodied in individual humans, and disembodied knowledge. He focuses on the properties of the latter category, the world of ideas and research, supposing that there is sufficient dynamic feedback in the research sector to generate endogenous growth and that the scope for developing new ideas is limitless.

In Romer's model, it is the number of people engaged in research and development that drives long-run growth. His mathematical representation of the generation of new ideas (or blueprints for new products) is similar to that of Lucas's educational sector:

$$\frac{dA_t}{dt} = \phi L_A A_t, \tag{6}$$

where A_t represents the number of productive ideas that have been realized at time t in history, and the differential, dA/dt, is the current output of new ideas from the research sector. L_A represents the amount of human capital, or the number of researchers, devoted to innovation.

Crucially, Romer assumes that the rate of innovation is directly proportional to the extant stock of knowledge. This is the "standing on shoulders" hypothesis of knowledge accumulation, so labelled by Charles Jones (1998), in reference to Isaac Newton's disclaimer: "If I have seen farther than others, it is because I was standing on the shoulders of giants."

In the accumulation of disembodied ideas, rather than embodied skills, it is indeed plausible to suppose that the level of current output might be directly proportional to the size of the stock. The more ideas and theorems that we have to draw on, the easier it is to generate new ones. Moreover, ideas do not necessarily disappear when their developer dies: They can typically be recorded and transmitted at minimal cost.

Implicit in Romer's formulation of research output is the idea that there is an evenly distributed and infinite universe of potential ideas waiting to be discovered. Thus, a given amount of research effort will produce a predictable number of new ideas. A more realistic approach, allowing the discovery rate to fluctuate, is summarized by Aghion and Howitt (1998) in their discussion of general-purpose technologies stemming from innovations such as the steam engine, the electric dynamo, and the computer.

4. Lucas (1988) asserts that his model of endogenous growth can be sustained across generations if a child's initial endowment of human capital is proportional to the level already attained by the adults—but, unless Lysenko was correct, the genetic transmission of acquired human capital is unlikely.

1.1.5 Nonrivalry of Ideas

As well as hypothesizing dynamic feedback in the generation of new ideas, Romer emphasizes that ideas have another significant economic property, *nonrivalry.* Objects are usually rival, meaning that if you are using something, I cannot use it at the same time. But this is not true of ideas. Once the binomial theorem has been published, your use of it does not in any way interfere with my use of it.

Of course, people can try to stop others from making use of patented ideas. But the *excludability* of ideas depends on the actions of people supported by institutions of laws and property rights, rather than nonrivalry, which is an inherent feature of ideas.

Romer makes use of this distinction by assuming that ideas are fully excludable in their application to the production of goods. For example, a researcher can acquire full patent protection for the design of a new drug; it can only be manufactured if royalties are paid. On the other hand, she has no protection against other researchers who can reverse engineer her ideas and come up with their own different but improved drug design. Indeed, when the original researcher files her patent, she has to describe her idea, thereby providing her rivals with a free input into their subsequent research.

Romer's hypothesis that ideas are nonrival and nonexcludable in the research process has important implications for public policy. Researchers may reap the benefits from the direct application of their ideas, but they do not receive monetary reward from others who "stand on their shoulders." Left to the market, there will be an undersupply of research effort. Public intervention is required to subsidize research, hence to stimulate growth, up to the socially optimum level.

Other aspects of knowledge accumulation are analysed by Aghion and Howitt (1992, 1998), who emphasize the Schumpeterian notion of "creative destruction." Patent rights may bestow monopoly power on the producer of a particular generation of an innovative good, but they cannot prevent the development by a rival of the next generation of goods that are superior in quality and/or price. The creation of the improved version destroys the flow of profits to the previous monopolist. Unbridled competition in such a market can lead to too much research being carried out, where the research is concerned with marginal quality improvements rather than new products and processes. Nevertheless, such research is still capable of driving long-run economic growth.

1.2 The Cambridge Counterrevolution

The intellectual euphoria of endogenous growth theory was challenged by a group of economists, mostly connected with or based in Cambridge, Massachusetts, who chose to stand behind (or on the shoulders of) Nobel laureate Robert Solow of the Massachusetts Institute of Technology.

Solow (1994) himself is critical of the knife-edge assumption required to generate stable long-run growth in the models of Romer and Lucas. His point is that these models require the dynamic feedback parameter in the education/research sector to be exactly equal to unity. If we look back to Lucas's model, where the growth rate of the economy is determined by equation (5), we can see that a value of 0.9 for the parameter γ will, eventually, reduce growth to zero: The final term of that equation has $h_t^{1-\gamma}$ in the denominator, which drives the term to zero as human capital, h_t, rises if γ is less than one. Stable long-run growth requires that the parameter be exactly one.

Romer (1994) has argued that this knife-edge property can be overcome in a more complex model. More damaging to the endogenous growth cause, however, has been the empirical work of another Cambridge (Harvard)–based economist, Greg Mankiw. In a much-cited paper—Mankiw, Romer, and Weil (1992)—he and his coauthors do not tackle the endogenous growth modelers head-on. Rather, they steal the ball of human capital from the endogenous growth scrum and use it to reconstruct the 1956 Solow model.

Their "augmented Solow model" includes human capital as a third factor in the aggregate production function, alongside capital and unskilled labor. They investigate the relationship between steady-state levels of output and the three inputs, using secondary-school enrollment rates as a proxy for the rate of investment in human capital. They conclude that the factors are of approximately equal importance—that is, that the elasticity of output with respect to each factor is approximately one-third—and that together they account for 80 percent of the observed variation in 1985 income levels across some ninety-eight nations.

This was a neat sidestep, rather than a direct hit on endogenous growth theory. There was no attempt to directly confront the two models with a discriminating statistical test, but the 1956 model was effectively rehabilitated—even though the econometric evidence is likely to be flawed due to the endogeneity of the explanatory variables. Moreover, this was only half-time in the comeback match. In an equally influential second half, the Mankiw, Romer, and Weil (1992) paper provided a clever reinterpretation of an empirical regularity. Studies of postwar economic growth had typically reported a conditional convergence effect. These studies ran regression models of the form

$$(7) \qquad (\ln y_{iT} - \ln y_{i0}) = \alpha_0 + \beta \ln y_{i0} + \mathbf{gX_i} + \varepsilon_i,$$

where the dependent variable is the growth rate of y, output per capita (or per worker), over a period of T years. $\mathbf{X_i}$ represents a vector of additional explanatory variables. "Conditional convergence" is said to exist if the regression parameter, β, is negative—a lower starting value for y is associated with a higher subsequent rate of growth, conditional on the **X** variables that explain differences in rates of growth.

Previous authors[5] had interpreted conditional convergence as evidence that technological spillovers from the most advanced economies enabled less advanced economies to imitate and thus enjoy relatively fast productivity growth. The Mankiw-Romer-Weil reinterpretation of such evidence, echoed by their Harvard colleagues Robert Barro and Jeffrey Sachs,[6] involves treating the **X** variables as determinants of the neoclassical steady state, rather than the long-run growth rate. They then interpret the initial income variable ($\ln y_0$) as a measure of distance from steady state and the β-coefficient as a measure of the speed of convergence to steady state.

This reinterpretation of the evidence in favor of the neoclassical model has been complemented by the more direct approach of MIT graduate Charles Jones.[7] He highlights the fact that endogenous growth models based on the accumulation of knowledge, such as Romer's model, typically suggest that the rate of growth should be an increasing function of the resources devoted to R&D.[8] He cites evidence from the United States that contradicts this prediction: "Since 1950, the fraction of the labour force engaged in formal R&D has increased by almost a factor of three. Despite these changes, average growth rates . . . are no higher today than they were from 1870 to 1929" Jones (1998, 157).

Jones also criticizes some of the key assumptions underpinning the knowledge-based models of endogenous growth. In particular, he suggests that knowledge creation may become more difficult over time as the easy ideas are discovered first, leaving subsequent researchers with a pool that has been "fished out." He also suggests that researchers may often duplicate each other's efforts, "stepping on toes" rather than "standing on shoulders."

These critiques of endogenous growth theory seem to imply that policies aimed at increasing investment in education and/or research will not be successful in raising the rate of economic growth for a sustained length of time. I will argue in the next section of the paper that this is not necessarily the case.

1.3 Reconciling Conflicting Theories of Growth

A crucial difference between the neoclassical and new growth theories concerns the question of whether the long-run rate of growth of the econ-

5. For example, Abramovitz (1986) and Dowrick and Nguyen (1989).
6. See Barro and Lee (1994) and Sachs and Warner (1997).
7. See Jones (1995a,b).
8. However, Aghion and Howitt (1998) show that their Schumpeterian model of endogenous innovation can be adapted to eliminate the scale effect.

omy is some exogenous constant or whether it can be influenced by public policy. Put another way, the question is whether policies and institutions that influence the rate of accumulation of physical and/or human capital have long-run effects on the *level* of economic activity or on its *rate of growth.* For purposes of practical policy making, however, this distinction may be relatively unimportant—if the long run never arrives. Looking back to figure 1.1, if economies are subject to shocks of sufficient magnitude and frequency, it may be difficult, if not impossible, to tell whether the long-run growth path really looks like path B or path C. In the short run—between time T_0 when the first major shock occurs and some time T_1 when another such event occurs—the paths may be virtually indistinguishable.

The evidence of the neoclassical revivalists can be interpreted to support this view. Mankiw, Romer, and Weil (1992), Barro and Sala-i-Martin (1995), and Sachs and Warner (1997) all report growth regression evidence suggesting that the rate of convergence toward steady state is of the order of 2 percentage points per year, implying that it will take more than thirty years for a country to halve the gap between its current income and the steady-state level.[9] Within a half-life of several decades, we must surely expect that there will be changes in investment rates and changes in the rate of technological progress such that the neoclassical economy is rarely able to get close to steady state.

A useful way to think of this problem is to consider the specification of the Error Correction Model (ECM). The ECM is commonly used to decompose macroeconomic time series into cyclical and long-run components and to test for long-run cointegrating relationships. A typical regression is of the form

$$(8) \qquad (\ln y_t - \ln y_{t-1}) = \alpha\{\mathbf{X}_t - \mathbf{X}_{t-1}\} + \lambda[\ln y_{t-1} - \beta\mathbf{Z}_{t-1}] + \varepsilon_t,$$

where y represents real output and the dependent variable is the growth rate of output. The explanatory variables are segregated. The **X** variables, which influence short-run movements, are entered in first differences. The **Z** factors are entered as lagged variables, along with the lagged value of output, y_{t-1}. For analysis of the long-run path, the first differences are set to zero, yielding the long-run path for output as a function of the **Z** variables:

$$(9) \qquad \ln y_{T^*} = \beta\mathbf{Z}_{T^*}$$

9. Subsequent studies on panel data have estimated higher speeds of convergence: In particular, Islam (1995) and Lee, Pesaran, and Smith (1997) estimate annual convergence rates up to 9 percent and 30 percent respectively. But Dowrick and Rogers (2002) show that these studies confound the effects of neoclassical convergence—due to diminishing returns to investment—with the effects of international technology diffusion. Separating out these effects, they find that the half-life of neoclassical convergence is more than fifteen years.

This very general empirical specification is consistent with both exogenous growth and endogenous growth models. If the **Z** vector contains a time trend, T, the regression coefficient on T is an estimate of the exogenous rate of technological progress—as in the neoclassical model. However, the **Z** vector may equally well contain the time trend interacted with another variable, such as the level of human capital. If so, the coefficient on this term captures the impact of human capital on the long-run growth rate of the economy—as predicted by some endogenous growth models.

In the ECM framework, the sign of the regression coefficient λ indicates whether output converges to the long-run path. The square brackets in equation (8) capture last period's deviation from the long-run path. The negative value of λ indicates the proportion of last period's "error" that is "corrected" in the current period.

A typical time series study that is trying to identify breaks in trend growth, using thirty to forty annual observations, might find a half-life for the business cycle of two to three years.[10] In this context, the "trend" growth is approximated by the average growth rate over one or two decades, averaging out fluctuations over three or four business cycles. But if convergence to the neoclassical steady-state growth path has a half-life of thirty years, this time scale is clearly insufficient to capture the underlying long-run rate of growth. Rather, we are identifying changes in the slope of the transitional growth path.

This supposition is confirmed by the recent study of Jones (2001). He adopts a modified growth accounting approach to analyze the last fifty years of U.S. growth. He finds that only one-fifth of the actual growth rate of labor productivity (averaging 2.0 percent per year) has been attributable to exogenous technical change. The remaining four-fifths of growth (1.6 percent per year) is attributable to continued growth in education and research intensity. In his terms, "Transition dynamics associated with educational attainment and the growth in research intensity account for 80 percent of growth" (p. 23).

Jones's conclusion is couched in the language of the neoclassical approach. Sustained growth above steady-state levels can only be transitional and is driven by sustained (but ultimately bounded) growth in the share of GDP going to investment in human capital. An alternative interpretation of the same evidence might claim that increased investment in human capital has raised the long-run endogenous rate of growth.

Evidence that reconciles the two approaches to understanding growth comes from Benhabib and Spiegel (1994), who carry out econometric esti-

10. A pooled time series cross-section study by Lee, Pesaran, and Smith (1997), allowing for heterogeneity in country-specific time trends, has estimated convergence in the Solow-Swan model to have a half-life of 2.5 years. I interpret this as a failure to distinguish the speed of transition to steady state from the fluctuations of the business cycle.

mation on various models to explain variation in twenty-year growth rates (1965–85) on a cross section of seventy-eight countries. In their preferred model, technological progress is the sum of two components: an exogenous component, as in the neoclassical model, and a semi-endogenous component, related to the rate of absorption of technology from the technological leading country, captured by an interactive term between the productivity gap and the level of human capital. Their preferred model draws on the analysis of Nelson and Phelps (1966).

They report that the interactive term is statistically significant, supporting the idea that there is an endogenous component to technological progress. At the same time, they estimate an output elasticity close to 0.5 for physical capital, suggesting diminishing returns to investment and a slow rate of convergence toward the steady-state capital stock.

Broadly similar results are reported by Dowrick and Rogers (2002). Our study differs from that of Benhabib and Spiegel (1994) in that we carry out the analysis on a panel of growth data. This enables us to test for country-specific effects. We also use an instrumental variable estimator to control for reverse causation between growth and the explanatory variables. Country-specific effects, which we interpret as endogenous components of technical progress, are found to be important. We confirm the finding that the level of human capital facilitates technological catch-up, especially among the middle-income and richer countries.

These models combine features of the neoclassical theory with the new growth theory. Changes in the rate of physical investment have, ultimately, only level effects; however, within a time frame of one or two decades this is indistinguishable from a growth effect. At the same time, countries have different rates of technological progress with an endogenous component, dependent on the stock of human capital and the allocation of resources to research, and a semi-endogenous component, dependent on the rate of technological change at the frontier and on the country's ability to absorb ideas from abroad.

1.4 Evidence on Education and Growth

Some of the earliest studies that investigated the link between education and economic growth were conducted by Mankiw, Romer, and Weil (1992) and Barro (1991). They examined variations in school enrollment rates, using a single cross section of both the industrialized and the less-developed countries. Both studies concluded that schooling has a significantly positive impact on the rate of growth of real GDP. They interpreted this as evidence of changes to (short-run) transitional growth paths. Barro and Sala-i-Martin (1995) also investigated the impact of educational expenditures by governments, finding that they have a strong positive impact. Using instrumental variable techniques to control for simultaneous causation, their regressions suggest that the annual rate of return on public education is of the order of 20 percent.[11]

A series of subsequent studies made use of panel data, examining changes over time in both education and growth. Several of these panel studies—including Benhabib and Spiegel (1994), Islam (1995) and Caselli, Esquivel, and Lefort (1996)—failed to detect any significant relationship between the rate of increase of educational capital and the rate of economic growth. They suggested that the positive findings of the earlier cross-section studies were due to omitted variable bias, failing to control for country-specific effects.

More recently, a third generation of studies has suggested a number of reasons why the negative findings of previous panel studies might have been biased. Pritchett (2001) has argued that poor policies and institutions have hampered growth in many of the least developed economies, directing skilled labor into relatively unproductive activities, hence disrupting the statistical relationship between education and growth in samples that include less-developed economies. Krueger and Lindahl (2001) suggest that the problem of unobserved variation in educational quality is exacerbated in panel data. Taking data quality into account, they show that increases in the stock of schooling do improve short-run economic growth. Hanushek and Kimko (2000) confirm that direct measures of labor force quality, from international mathematics and science test scores, are strongly related to growth. Temple (2001) finds that growth effects are positive but nonlinear. These nonlinear effects may be missed by studies that impose linearity.

Overall, it seems that studies that pool the least and the most developed economies do not find consistent and robust relationships between education and growth. For evaluation of Australian policy, it is probably more useful to examine studies that are restricted to Organization for Economic Cooperation and Development (OECD) economies.

Mankiw, Romer, and Weil (1992) estimate the determinants of countries' steady-state income levels as a function of investment in both physical and human capital. For their cross section of OECD countries, they estimate an elasticity of 0.76 between steady-state output and the proportion

11. Barro and Sala-i-Martin (1995) report an increasing marginal effect on growth of years of schooling, but this may be due to a lack of variation in the data on primary enrollments. More surprising is their finding that positive growth effects are confined to male education. On the other hand, a study by Caselli, Esquivel, and Lefort (1996) uses a more sophisticated panel estimation technique (general method of moments) and reverses the result—it is female secondary education rather than male education that promotes growth. This finding is confirmed by Knowles, Lorgelly, and Owen (2002). These contradictory results probably reflect strong colinearity between female education, male education, and other measures of development, such as life expectancy and fertility, which are included in the regressions. Moreover, where many women are involved in domestic rather than market economic activity, the educational enhancement of their contribution to economic welfare may not be picked up directly by standard measures of GDP.

Table 1.1 Predicted Increase in the Level of Output for an Additional Year of Schooling in the Adult Population of an OECD Country

Study	Level Effect (%)
Bassanini and Scarpetta (2002)	6
Mankiw, Romer, and Weil (1992)	6–15

of the workforce enrolled in secondary school. Translating the elasticity into the marginal impact of an additional year of schooling in OECD countries (where average schooling varies between five and twelve years), this implies that steady-state real GDP increases in a range of 6 to 15 percent, with an estimated 8 percent increase for a country like Australia with average schooling of ten years.

Bassanini and Scarpetta (2002) analyze panel data, using annual data for twenty-one OECD countries from 1971 to 1998. They use a pooled mean group estimator, which allows for cross-country variations in short-run coefficients, but they test for and impose homogeneity on long-run coefficients. Their most reliable estimates suggest that the return to an additional year of schooling is a 6 percent increase in steady-state output. Table 1.1 summarizes these results.

These macroeconomic estimates refer to that part of the social returns to schooling that is captured in GDP. It appears that these estimates are close in magnitude to microeconomic estimates of private returns to the education of individuals. This implies that the external effects of education are relatively small, at least in the context of the *level* effects of education.

These conclusions must be modified, however, in the light of a series of empirical studies that have been inspired by the hypothesis of Nelson and Phelps (1966) that human capital may influence the rate of introduction of new technologies. Benhabib and Spiegel (1994), for example, compare models that treat human capital as a direct input into production with models treating human capital as an intermediate input into the acquisition of skills and/or knowledge. The former implies a relationship between output growth and educational growth, whereas the latter implies a relationship between output growth and the average stock of human capital per worker. Their econometric evidence favours the latter model. A more educated workforce can more readily identify, adapt, and implement new ideas—whether the ideas are generated domestically or overseas.

This finding, that education levels affect long-run technological progress, is confirmed by Frantzen (2000), who analyzes the growth of total factor productivity (TFP) between 1961 and 1991 in the business sectors of twenty-one OECD countries. It is also confirmed by Dowrick and Rogers (2002), who investigate the rate of technological convergence between 1970 and 1990 for a wide sample of fifty-one countries and for a sample of thirty-five relatively rich countries.

Table 1.2 Predicted Increase in Long-Run Economic Growth in Australia Due to an Additional Year of Schooling in the Adult Population

Study	Growth Effect (percentage points)
Benhabib and Spiegel (1994)	0.3
Frantzen (2000)	0.8
Dowrick and Rogers (2002)	0.2–0.5[a]

[a]The lower of these estimates is derived using the coefficient reported in table 2 in Dowrick and Rogers (2002) using the full sample of countries. The higher estimate is from the coefficients in table 3, using the thirty-five-country sample of relatively rich economies with better data quality.

These studies share a common regression specification of the general form

$$\text{(10)} \qquad \text{TFP growth in country } i = \alpha S_i + \beta S_i f(\text{pr}_i) + \ldots,$$

where S_i is the average years of schooling in the adult population, and pr_i is the ratio of productivity in the technologically leading country relative to country *i*.

The first regression coefficient, α, captures the impact of schooling on domestic innovation. The second regression coefficient, β, captures the impact of schooling on the absorption of technological spillovers from the technologically leading country.

All three studies find that the level of schooling is a statistically significant determinant of growth. The predicted effect of an additional year of schooling in the adult population on the annual rate of growth of TFP is $\alpha + \beta f(\text{pr}_i)$. Considering the case of Australia, where the U.S. productivity ratio is approximately 1.5, we compare the predicted growth effects of schooling in table 1.2.

Even the lowest of these estimates predicts a highly significant boost to annual economic growth, one-fifth of a percentage point, for every additional year of schooling.

1.4.1 Australia's Educational Attainment Report

In the light of these estimates, it is of interest to draw up a report card on Australia's record of educational attainment. The data we use are taken from Barro and Lee (2001), who have revised and updated their previous estimates of the average years of schooling in the population aged twenty-five and over. Figure 1.2 shows the time path of this measure for Australia and selected OECD countries.

Forty years ago, Australian adults averaged 9.4 years of schooling, a level of attainment that not only was significantly above that of the other countries illustrated, but was surpassed only by New Zealand out of the 100 countries covered by Barro and Lee. By the year 2000, Australia's av-

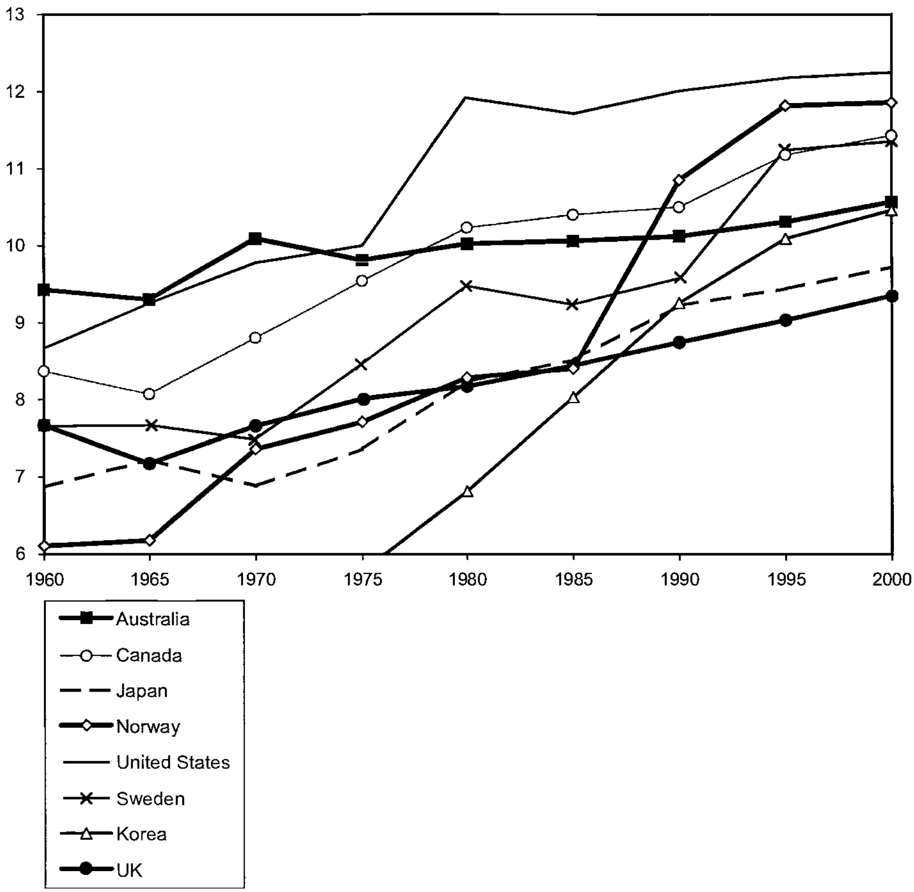

Fig. 1.2 Average years of schooling in the adult population
Source: Barro and Lee (2001).

erage schooling level had climbed to 10.6 years. Attainment rose faster, however, in all of our comparator countries, with the result that Australia has slipped below the United States, Norway, Sweden, and Canada, is only fractionally higher than Korea, and is only slightly higher than Japan.

Of course, the average of years of schooling is an imperfect measure of skills and abilities, since educational quality varies across countries and over time, and because it ignores the abilities acquired through experience and workplace training. In the mid-1990s, twenty countries participated in the OECD's International Adult Literacy Survey. This survey provides a direct comparison of work-related skills, including measures of literacy and numeracy. Figure 1.3 presents a scatter plot that demonstrates that, on either measure, Australian adults rate close to the OECD average.

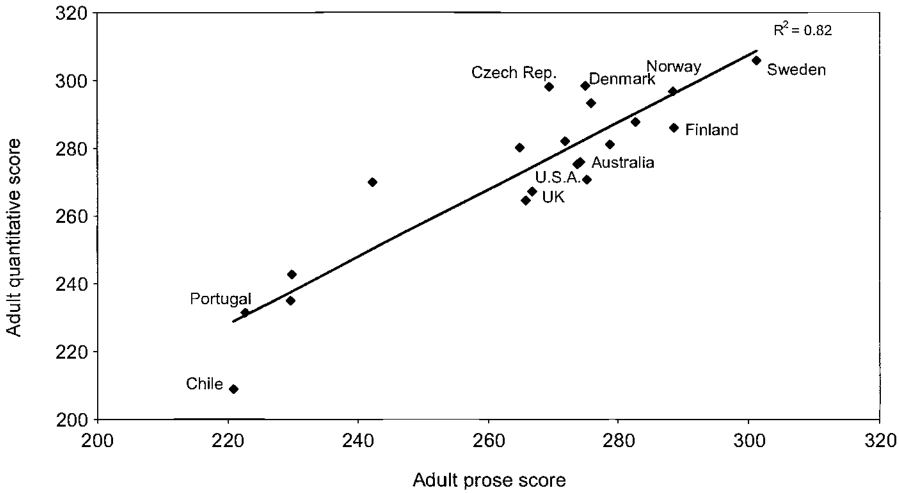

Fig. 1.3 Quantitative and verbal skills
Source: Barro and Lee (2001).

The Third International Maths and Science Study, conducted in 1994 and 1995, confirms Australia's average performance. On measures of seventh-grade proficiency in math and science, Australian school students ranked fifteenth and twelfth respectively out of the thirty-seven country scores reported by Barro and Lee (2001).

These international comparisons suggest that Australia's educational report card should be marked "Started well, but slacked off. Substantial room for improvement."

1.5 The Contribution of Research and Development to Economic Growth

I have already discussed the attributes of knowledge that make it significantly different from the accumulation of items of physical capital. These special attributes are nonrivalry and dynamic feedback. Once a new idea has been generated, it can be used simultaneously and costlessly in many different processes. Furthermore, the idea can serve as an example and inspiration for further research.

These are the attributes of knowledge that give it the potential to drive long-run growth. But the properties of nonrivalry and feedback also suggest that the market may fail to allocate sufficient resources to knowledge generation because individuals have difficulty in establishing and enforcing property rights over their new ideas: Some of the benefits of an innovation are likely to accrue to others. When the private return to innovation is less than the social return, governments need to subsidize R&D.

Paul M. Romer (1993) has argued that while governments should fund fundamental research, it may well be appropriate for self-funding industry associations to fund development and applied research—using governments only to enforce the collection of agreed contributions.[12] Weder and Grubel (1993) expand this point in their discussion of the "Coasean" institutions that operate in various countries to internalize knowledge spillovers and promote technical progress. In particular, they cite the occurrence of three types: (1) industry associations such as the Japanese *keiretsu* or Swiss Verbande; (2) conglomerate corporations, including multinational enterprises; and (3) geographic clustering of industries, such as Silicon Valley or the Northern Italian networks. They point particularly to the Swiss and Japanese examples, where voluntary associations, supported by public policy, encourage long-run relationships between vertically related firms and encourage joint ventures and cooperation including joint research and training schemes.

Expenditures on R&D typically constitute, for advanced economies, only a few percent of GDP—perhaps one-tenth of the expenditure devoted to investment in physical equipment and structures. In a standard growth accounting framework, variations in research effort will, therefore, explain very little of the differences in growth rates between countries. But the point of much of the new growth theory is precisely that if knowledge spillovers are substantial, and if knowledge exhibits dynamic feedback effects, then even small changes in the resources devoted to the production of knowledge may result in substantial changes in economic growth. This point is made by Grossman and Helpman (1991), who calibrate their model to match the U.S. growth experience. They predict that, while business investment constitutes around 10 percent of GDP, investment in R&D—the engine of growth—need comprise as little as 1.6 percent to generate economic growth of 2.5 percent per year.

Lichtenberg and Siegel (1991) surveyed some fifteen previous studies into R&D investment by U.S. firms and industries, reporting real private rates of return averaging 25 percent. Their own econometric study of two thousand U.S. firms revealed a 30 percent rate of return on company-funded R&D, a "productivity premium" on basic research, and a 7 percent return on federally funded company research. These estimates of private rates of return on company-funded R&D are very high, given that investment in physical capital might be expected to earn a return closer to 10 percent. The higher rate of return on research reflects, presumably, a large premium for risk and problems in diversifying or pooling such risk.

Nadiri (1993) confirms that private returns to R&D are particularly high in his review of the literature: "[N]et rates of return on own R&D of 20% to 30% at the firm level and 10% to 30% at the industry level are reasonable sets of estimates." He goes on to examine spillovers to other industries and

12. Australian agricultural research has long been funded on this basis.

concludes that "The spillover effects of R&D are often much larger than the effect of own R&D at the industry level. . . . [S]ocial rates of return often vary from 20% to over 100% with an average somewhere close to 50%" (pp. 34–35).

A subsequent paper by Lichtenberg (1992) is one of the first attempts at studying the cross-country evidence on the impact of R&D expenditures on both the level and the rate of growth of real GDP. Using a sample of seventy-four countries, his growth regressions, using the neoclassical framework, reveal that returns to R&D are approximately double the returns to physical investment—a result that is broadly consistent with estimates from the microeconomic studies of firms and industries.

Coe and Helpman (1995) try to quantify the magnitude of international R&D spillovers. They seek to explain variations in the annual growth of TFP for twenty-one OECD countries, plus Israel, over the period 1970–90. Their econometric analysis finds that the stock of knowledge in one country, measured by cumulated historical R&D expenditures, raises productivity in foreign countries with which they trade. It is not clear exactly why the extent of technology transfer should depend on the magnitude of trade with a technologically advanced economy, although their empirical findings appear to be quite robust and have been confirmed by subsequent studies. One plausible explanation stems from the observation by Eaton and Kortum (2001) that the high R&D economies are also the major world exporters of capital goods. The general trade variable used by Coe and Helpman may be acting as a proxy for the import of high-tech capital goods for which the producers are unable to expropriate all of the rents.

Frantzen (2000) has extended the Coe and Helpman approach and provides us with estimates of rates of return on domestic R&D as well as estimating the strength of international technological spillovers. He finds that the following regression has strong statistical significance on a sample of twenty-one OECD countries:

The annual growth rate of TFP in the business sector, $1961-1991 = 0.59$

$\times$ (gross expenditure on own R&D)/(business-sector GDP) + 1.52

$\times$ SUM[(research intensity in country i) $\cdot$ (import share from country i)]

The first regression coefficient is an estimate of the national (social) rate of return to R&D—capturing not only the productivity benefits that accrue to the firms which make the investments but also the spillover benefits that accrue to firms in the same or related industries. The second regression coefficient captures the spillover benefit that a country can gain from research carried out by a trading partner. This benefit is proportional to the share of imports from that country in GDP—perhaps reflecting the embodied technological improvements in imported capital equipment.

It is instructive to compare Frantzen's estimates with other estimates of

Table 1.3 Estimated Rates of Return on R&D Expenditures

Study and Sample	Rates of Return (%) Private Returns	Social Returns	Cross-Country Spillovers
Lichtenberg and Siegel (1991)			
Survey of fifteen previous studies of U.S. firms and industries	25		
2000 U.S. firms	30		
Nadiri (1993)			
Survey of fifty U.S. and other studies at firm and industry level	20–30	50	
Lichtenberg and van Pottelsberghe de la Potterie (1996)			
GDP growth across OECD countries		51–63	
Frantzen (2000)			
Business-sector TFP growth across OECD countries		59	45[a]

[a]if the imports-GDP ratio equals 0.3.

returns to R&D expenditures. A summary is provided in table 1.3. Interestingly, he estimates a 59 percent social rate of return on national R&D expenditures, which is close to the average figure suggested by Nadiri's review of firm- and industry-level studies. It is also close to the results of Lichtenberg and van Pottelsberghe de la Potterie (1996), who estimate that the social rate of return on domestic R&D is 51 percent in the large Group of Seven (G7) economies but 63 percent in six smaller European countries.

All of these estimates lie substantially above the various estimates of private rates of return, implying that there are very significant spillover effects between the firms and industries within a national economy.

The implication for Australia of the benchmark Frantzen estimate can be calculated as follows. Our gross annual R&D expenditure (public and private combined) of around ten billion dollars amounts to 1.5 percent of total GDP, or approximately 2 percent of business-sector value added. An additional billion dollars' annual expenditure on R&D, representing one-fifth of 1 percent of value added, is predicted to increase the annual growth rate by just over one-tenth of a percentage point.

What would happen if the countries from which Australia imports capital goods were each to increase their research intensity by 0.2 percentage points (the same rise as in the example for Australia)? If we multiply the regression coefficient on foreign R&D by Australia's total share of imports in GDP, which is 30 percent, we find that technology spillovers are predicted to increase Australian growth by just over one-tenth of a percentage point. In other words, domestic R&D and spillovers from foreign R&D are of roughly equal importance for productivity growth.

1.5.1 Australia's Investment in Research and Development

Compared with the leading industrial economies of the OECD, Australia invests less of its resources into R&D—and a lesser proportion of that investment is carried out within the business sector. In figure 1.4 we see that the share of GDP devoted to R&D in Australia has been growing over the past few decades, from under 1 percent to around 1.5 percent. Our research intensity is, however, still well below that achieved by major industrial economies such as Japan, the United States, and Germany, where the R&D ratio has averaged 2.5 percent over the past twenty years. On the other hand, Australian R&D intensity is close to or above that of Canada and New Zealand, countries with comparably large rural sectors. R&D intensity dipped after peaking at 1.7 percent in 1996. This is attributable in the first instance to a fall in R&D within the business sector of the Australian economy, which was driven in part by the reduction in the tax concession for R&D.

Even at its peak, the Australian business sector's contribution to R&D has been comparatively low. In figure 1.5 we see that the proportion of total R&D that is carried out in the business sector had been rising from 1981 up to 1995, from 25 percent to 51 percent, but then fell to 45 percent by 1998. In the economies that are illustrated in figure 1.6, with the exception of Australia and New Zealand, well over half of national R&D was carried out by the business sector.

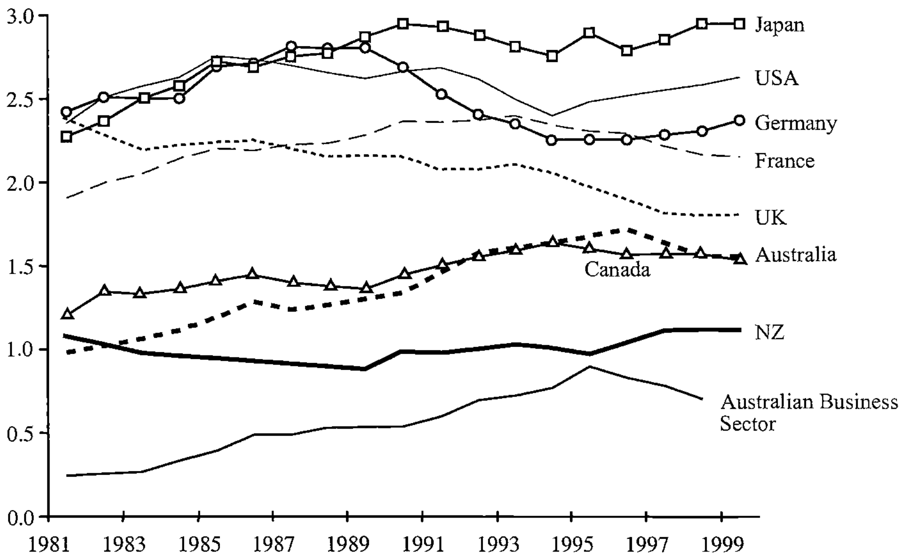

Fig. 1.4 Gross expenditure on R&D as % of GDP, 1981–99
Source: OECD R&D database.

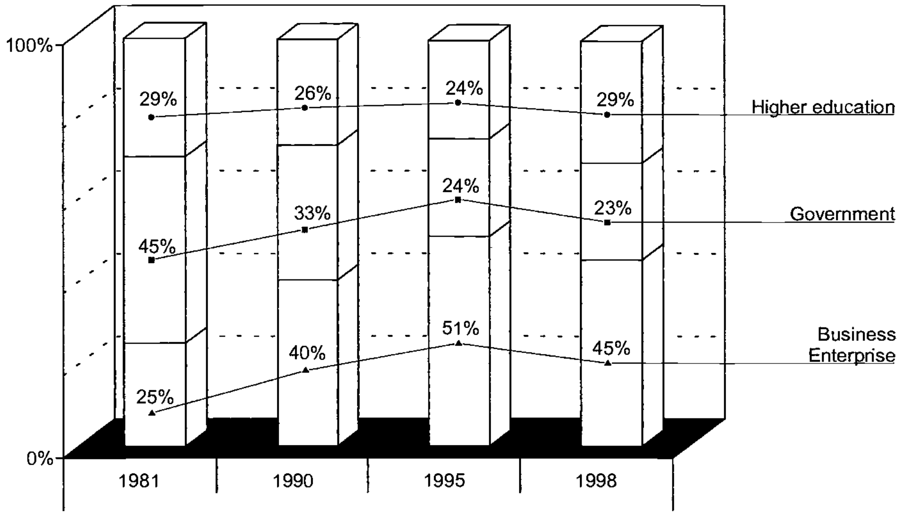

Fig. 1.5 Sectoral composition of Australian R&D
Note: The residual category is the share of private nonprofit R&D.

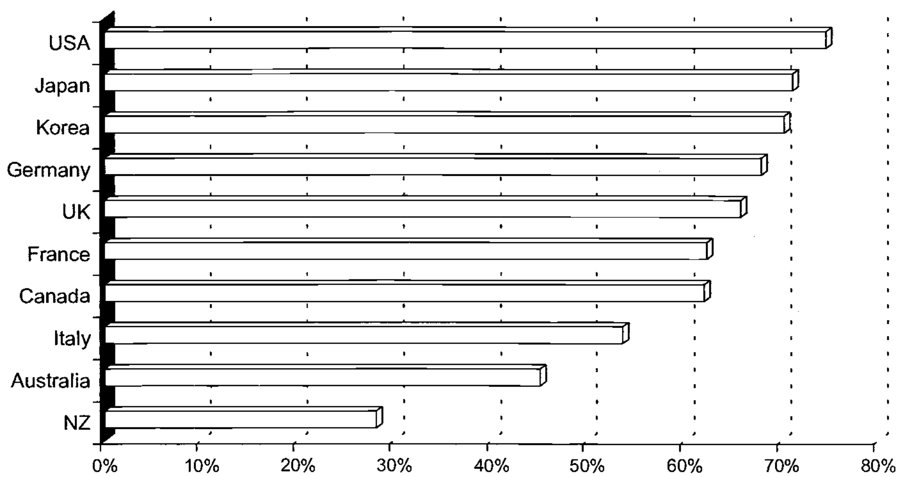

Fig. 1.6 Business enterprise R&D as a share of total (1988)
Source: OECD R&D database.

An interesting perspective on Australian performance on a broader measure of "investment in knowledge" comes from OECD (2001), which aggregates expenditures on R&D, higher education, and computer software. On this measure, Australia ranks fourteenth out of the twenty-four countries surveyed. In terms of the rate of growth of knowledge investment over the 1990s, Australia ranks tenth.

1.6 Conclusions

The neoclassical revival in growth theory has had the paradoxical effect of reinforcing one of the major points of the endogenous growth revolution. The driving force of economic growth is investment in human capital—skills and ideas—rather than investment in machines and buildings. The academic debate will no doubt continue over whether government policies that affect the rate of investment have any influence on the long-run, as well as the short-run, rate of growth of the economy. For practical purposes, however, if the "short run" involves a transition period of several decades, this debate may be strictly academic—in the pejorative sense of the word. Policies that affect investment, particularly in embodied or disembodied human capital, can have a sustained impact on economic growth.

A review of empirical studies on sources of economic growth confirms these claims: Both education and R&D are important sources of growth. In the mid-1990s, a number of studies were published claiming that there was no systematic relationship between changes in national educational attainment and changes in economic growth. Subsequent studies have, however, established that this lack of correlation was due to a mix of factors: poor institutional performance in some less-developed economies, and a failure to account for international variation in educational quality. Once we account for these factors, the evidence suggests returns to education that are consistent with microeconomic evidence on individual earnings. An increase of one year of schooling in the average educational attainment in the workforce, for example, can be expected to increase the long-run level of output by around 8 percent in a typical OECD country.

These are estimates of the *level* effects of education. A one-off increase in attainment will produce a one-off rise (albeit spread over time) in the level of GDP per capita. There is mounting evidence, however, that there are also substantial dynamic or *growth* effects, which are linked to a country's ability to implement new technologies. This evidence suggests that Australia would do well to increase its educational levels to match the OECD leaders: the United States and Scandinavia.

One of the concerns of current public debate is that the aging of the Australian population over the next fifty years will overtax (literally and metaphorically) the working-age population. From the perspective of growth theory, however, there may not be so much to fear. The aging of the demographic structure is being driven by the revolution in female education and workforce opportunities. For the generation born in the 1930s, only one-third of girls and one-half of boys completed high school. For the current generation, over 70 percent of boys and close to 80 percent of girls are completing year twelve. The past fifty years have also witnessed the end of legally enforced discrimination against women in the workforce—in the form of the marriage bar and legalized wage discrimination. These huge improvements in female education and workforce opportunities have been major factors in the fall in fertility, which is the driving force behind the changing age structure of the population.

The very factor that is causing the aging of the population, the revolution in women's education, gives us reason to expect continued strong growth of the Australian economy. The average educational attainment of the workforce will continue to rise for the next three decades as historical increases in school enrollments work their way through the adult population. These effects will be enhanced should educational enrollment continue to rise—particularly if the educational participation and achievement of Australia's young men rises to meet the levels of young women.

The evidence on the benefits of innovation is clear. A wide range of studies finds that private rates of return on R&D expenditures are very high, and that social rates of return—taking account of intranational spillovers of knowledge—are even higher.

We can summarize the potential productivity benefits for Australia of increased investment in education and research by using relatively conservative benchmark estimates, based on the large number of studies that have been summarized in this paper. Taking education first, an increase of 0.8 in the average years of schooling of the labor force would take us to 11.4 years, the average of the levels of attainment in North America and Scandinavia. The effect on the Australian economy would be an increase of one-third of a percentage point in the annual growth rate—coming both through human capital deepening and more rapid adoption of new technologies.[13]

Turning to investment in R&D, it is probably unrealistic to suppose that Australia will match the research intensity of the world leaders such as the United States or Germany. Adopting a more realistic role model, France, would require that an extra 0.6 percent of GDP be devoted to R&D—taking research intensity from 1.6 percent to 2.2 percent. Using a conservative estimate of the social rate of return,[14] the impact on the Australian economy would be an increase of one-quarter of a percentage point in the annual rate of productivity growth.

To sum up, positive prospects for continuing strong productivity growth will be enhanced if Australia emulates the higher rates of investment in knowledge—both in education and in R&D—that we observe in the leading OECD economies. An increasingly well-educated (albeit shrinking)

13. From table 1.1, the conservative estimate of the level effect is 0.8×6 percent $= 0.048$, which is equivalent to 0.0012 per year over forty years. From table 1.2, a conservative estimate of dynamic effect is $0.8 \times 0.003 = 0.0024$ per year. The two effects sum to 0.0036, or 0.36 percentage points per year.

14. Assuming the social rate of return is 0.4, which is substantially below the estimates summarized in table 3.

workforce, operating in an economy that continues to be open to trade in goods and ideas, will be well placed to identify, introduce, and manage the new technologies that will emerge over the next few decades.

References

Abramovitz, Moses. 1986. Catching up, forging ahead, and falling behind. *Journal of Economic History* 46:385–406.

Aghion, Philippe, and Peter Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60 (March): 323–51.

———. 1998. *Endogenous growth theory.* Cambridge, Mass.: MIT Press.

Apps, Patricia, and Ray Rees. 2001. Household production, full consumption and the costs of children. *Labour Economics* 8 (6): 621–48.

Ashenfelter, Orley, and Alan B. Krueger. 1994. Estimates of the economic returns to schooling from a new sample of twins. *American Economic Review* 84 (5): 1157–73.

Ashenfelter, Orley, and Cecilia Rouse. 1997. Income, schooling, and ability: Evidence from a new sample of identical twins. NBER Working Paper no. 5688. Cambridge, Mass.: National Bureau of Economic Research, July.

———. 1998. Income, schooling, and ability: Evidence from a new sample of identical twins. *Quarterly Journal of Economics* 113 (1): 253–84.

Barro, Robert J. 1991. Economic growth in a cross section of countries. *Quarterly Journal of Economics* 106 (2): 407–43.

Barro, Robert J., and Jong-Wha Lee. 1994. Sources of economic growth. *Carnegie Rochester Conference Series on Public Policy* 40 (0): 1–46.

———. 2001. International data on educational attainment: Updates and implications. *Oxford Economic Papers* 53 (3): 541–63.

Barro, Robert J., and Xavier Sala-i-Martin. 1995. *Economic growth.* New York: McGraw-Hill.

Bassanini, Andrea, and Stefano Scarpetta. 2002. Does human capital matter for growth in OECD countries? A pooled mean-group approach. *Economics Letters* 74:399–405.

Benhabib, Jess, and Mark Spiegel. 1994. The role of human capital in economic development: Evidence from aggregate cross-country data. *Journal of Monetary Economics* 34 (2): 143–73.

Borland, Jeff, and Xiaokai Yang. 1992. Specialization and a new approach to economic organization and growth. *American Economic Review* 82 (2): 386–91.

Caselli, Francesco, Gerardo Esquivel, and Fernando Lefort. 1996. Reopening the convergence debate: A new look at cross-country growth empirics. *Journal of Economic Growth* 1 (September): 363–89.

Coe, David T., and Elhanan Helpman. 1995. International R&D spillovers. *European Economic Review* 39 (5): 859–87.

Dowrick, Steve, and Duc Tho Nguyen. 1989. OECD comparative economic growth 1950–85: Catch-up and convergence. *American Economic Review* 79 (5): 1010–30.

Dowrick, Steve, and Mark Rogers. 2002. Classical and technological convergence: Beyond the Solow-Swan growth model. *Oxford Economic Papers* 54:369–385.

Eaton, Jonathan, and Samuel Kortum. 2001. Trade in capital goods. NBER Work-

ing Paper no. 8070. Cambridge, Mass.: National Bureau of Economic Research, January.
Folbre, Nancy, and Julie A. Nelson. 2000. For love or money—or both? *Journal of Economic Perspectives* 14 (4): 123–40.
Frantzen, Dirk. 2000. R&D, human capital and international technology spillovers: A cross-country analysis. *Scandinavian Journal of Economics* 102 (1): 57–75.
Grossman, Gene M., and Elhanan Helpman. 1991. Trade, knowledge spillovers, and growth. *European Economic Review* 35:517–26.
Hanushek, Eric A., and Dennis D. Kimko. 2000. Schooling, labor-force quality, and the growth of nations. *American Economic Review* 90 (5): 1184–208.
Islam, Nazrul. 1995. Growth empirics: A panel data approach. *The Quarterly Journal of Economics* 110 (443): 1127–70.
Jones, Charles I. 1995a. R&D-based models of economic growth. *Journal of Political Economy* 103 (4): 759–84.
———. 1995b. Time series tests of endogenous growth models. *Quarterly Journal of Economics* 110 (2): 495–525.
———. 1998. *Introduction to economic growth.* New York: W. W. Norton.
———. 2001. *Sources of U.S. economic growth in a world of ideas. American Economic Review* 92 (1): 220–239.
Jones, Charles I., and John C. Williams. 2000. Too much of a good thing? The economics of investment in R&D. *Journal of Economic Growth* 5 (1): 65–85.
Knowles, Stephen, Paula K. Lorgelly, and P. Dorian Owen. 2002. Are educational gender gaps a brake on economic development? Some cross-country empirical evidence. *Oxford Economic Papers* 54:118–49.
Krueger, Alan B., and Mikael Lindahl. 2001. Education for growth: Why and for whom? *Journal of Economic Literature* 39:1101–36.
Lee, Kevin, M. Hashem Pesaran, and Ron Smith. 1997. Growth and convergence in a multi-country empirical stochastic Solow model. *Journal of Applied Econometrics* 12:357–92.
Lichtenberg, Frank R. 1992. R&D investment and international productivity differences. NBER Working Paper no. 4161. Cambridge, Mass.: National Bureau of Economic Research.
Lichtenberg, Frank R., and Donald Siegel. 1991. The impact of R&D investment on productivity: New evidence using linked R&D-LRD data. *Economic Inquiry* 29 (April): 203–28.
Lichtenberg, Frank, and Bruno van Pottelsberghe de la Potterie. 1996. International R&D spillovers: A re-examination. NBER Working Paper no. 5668. Cambridge, Mass.: National Bureau of Economic Research, July.
Lucas, Robert E., Jr. 1988. On the mechanics of economic development. *Journal of Monetary Economics* 22 (1): 3–42.
Mankiw, N. Gregory, David Romer, and David N. Weil. 1992. A contribution to the empirics of economic growth. *Quarterly Journal of Economics* 107 (2): 407–37.
Miller, Paul W., Charles Mulvey, and Nick Martin. 1995. What do twins studies reveal about the economic returns to education? A comparison of Australian and U.S. findings. *American Economic Review* 85 (3): 586–99.
Nadiri, M. Ishaq. 1993. Innovations and technological spillovers. NBER Working Paper no. 4423. Cambridge, Mass.: National Bureau of Economic Research.
Nelson, Richard R., and Edmund S. Phelps. 1966. Investment in humans, technological diffusion, and economic growth. *The American Economic Review* 56 (1/2): 69–75.
Organization for Economic Cooperation and Development (OECD). 2001. *STI Scoreboard.* Paris: Author.

Pitchford, John. 1960. Growth and the elasticity of factor substitution. *Economic Record* 36 (December): 491–504.
Preston, Alison. 1997. Where are we now with human capital theory in Australia? *The Economic Record* 73 (220): 51–78.
Rebelo, Sergio. 1991. Long-run policy analysis and long-run growth. *Journal of Political Economy* 99 (3): 500–21.
Romer, Paul. 1993. Idea gaps and object gaps in economic development. *Journal of Monetary Economics* 32 (3): 543–73.
Romer, Paul M. 1986. Increasing returns and long-run growth. *Journal of Political Economy* 94 (5): 1002–37.
———. 1990. Endogenous technological change. *Journal of Political Economy* 98 (5): S71–S102.
———. 1993. Implementing a national technology strategy with self-organizing industry investment boards. *Brookings Papers on Economic Activity, Microeconomics:* 345–90.
———. 1994. The origins of endogenous growth. *Journal of Economic Perspectives* 8 (1): 3–22.
Sachs, Jeffrey D., and Andrew M. Warner. 1997. Fundamental sources of long-run growth. *American Economic Review* 87 (2): 184–88.
Shi, Heling L., and Xiaokai K. Yang. 1995. A new theory of industrialization. *Journal of Comparative Economics* 20 (2): 171–89.
Solow, Robert M. 1994. Perspectives on growth theory. *Journal of Economic Perspectives* 8 (1): 45–54.
Temple, Jonathan R. W. 2001. Generalizations that aren't? Evidence on education and growth. *European Economic Review* 45 (4–6): 905–18.
Weder, Rolf, and Herbert G. Grubel. 1993. The new growth theory and Coasean economics: Institutions to capture externalities. *Weltwirtschaftliches Archiv:* 488–513.
Yang, Xiaokai, and Jeff Borland. 1991. A microeconomic mechanism for economic growth. *Journal of Political Economy* 99 (31): 460–82.

Comment John Leahy

This paper provides a nice survey of the literature on knowledge and growth. At the heart of the discussion is the contrasting role played by knowledge in traditional growth theory and in the "new" growth theory.

In the traditional theory, as reformulated in the 1990s, knowledge is embodied in human capital. As its name suggests, human capital is just another form of capital. Human capital, like physical capital, is accumulated through investment, in this case investment in education. Like physical capital, it is subject to diminishing returns; given the supplies of the other factors, each additional unit of education adds a bit less to aggregate production. Modeling knowledge in this way does very little to alter the im-

John Leahy is professor of economics at New York University and a research associate of the National Bureau of Economic Research.

plications of traditional growth models, namely that accumulation of factors is subject to diminishing returns and does little to raise growth in the long run.

New growth theory makes two challenges to the traditional theory. First, by assuming constant returns to accumulable factors, it challenges the assumption that growth is exogenous in the long run. Second, by reinterpreting knowledge as ideas rather than human capital, it opens the door for the modeling of a number of interesting externalities.

While much has been made of the first challenge, I agree with Dowrick that the important contribution of the new growth literature lies in the reinterpretation of the role of knowledge and the novel externalities that arise. Endogenous growth depends on functional form assumptions that are convenient mathematically but for which we have no empirical evidence. We have no idea what would happen if suddenly the world found itself with twice the physical capital and twice the human capital, since this event has not yet happened.

The interpretation of knowledge as ideas has several implications that are relevant for public policy. First, ideas, unlike human capital, are nonrival. Absent legal restrictions, the fact that one person has an idea does not prevent others from having or using the same idea. The social returns to developing a new idea may therefore exceed the private return. The optimal government response may be to subsidize research. On the other hand, new ways of doing things tend to replace old ways of doing things. This creative destruction may mean that social returns are actually lower than private returns. Innovators do not take into consideration the losses to others that their innovations create. In such a case, the optimal government response would be to tax innovation.

The important question for public policymakers regards the size of these effects. The paper does a good job of surveying the empirical literature. Current studies tend to find large spillovers. Education in one country appears to raise productivity at home and abroad. The returns to R&D appear great, but the social returns appear greater still. These results appear to justify large subsidies to education and research and development.

At this point, I believe that the author could have been a bit more skeptical. I have severe doubts concerning both the magnitude and the interpretation of these results. As concerns interpretation, it is very difficult to identify a causal link between education and growth, let alone external effects of education within a country or the spillovers onto other countries. Not only may causality be reversed—with growth providing the resources for greater education—but countries with the foresight to educate their population probably get a lot of other things right as well. They may be more stable or have lower discount rates or better tax policy. They may subsidize other activities related to knowledge creation. They may have social systems that reward effort at all levels. It is very difficult to identify the effect of an exogenous shift in human capital induced by a government subsidy to education. Most studies do not even attempt to use instrumental variables, probably because there are few valid instruments.

The other problem is that some of the estimated effects are implausibly large. Most of the R&D studies cited in the paper find private rates of return in the neighborhood of 25 percent. This raises the question "Why are private agents not doing more R&D?" Firms appear to be leaving a lot of money on the table. Maybe, however, these returns are overstated and government-subsidized or -funded projects would yield substantially less. One possibility is that these estimates capture average rather than marginal rates of return. For example, economic research is valuable, but would doubling the money spent on economic research double the output of economic knowledge? Probably not. Another possibility is survivorship bias. We may not measure correctly all of the money spent on projects that don't pan out. Firms that fail are often dropped from the sample. In the end, the 7 percent return cited on federally funded R&D makes one wonder if policies would yield such amazing results.

In my mind, generating convincing estimates of the size of the externalities emphasized by the new growth literature remains one of the more important tasks facing growth economics.

Comment Andrew K. Rose

Steve Dowrick has given us a long, thorough, but focused survey of education and growth. He gives a fine summary of the theoretical literature that has obsessed much of the macroeconomics profession for the 1990s and concludes that the issue of whether growth is better modeled as being endogenous or exogenous may not actually be that relevant in practice. More important to my mind, he has also provided a number of estimates from both the micro- and macroeconomic literatures on the effects of education on output, and he concludes that they are large, even for an advanced OECD country like Australia. He believes there are large externalities and that the case for government intervention is secure. I agree with most of what he says, which seems reasonable in both the small and the large.

My personal view is that a survey like this should always focus precisely on a well-defined question. In this case, Dowrick is interested in answers to the question "What is the return to an additional year of education?" This creates a convenient taxonomy to organize the empirical estimates from

Andrew K. Rose is B.T. Rocca Professor of Economic Analysis and Policy at the Haas School of Business, University of California–Berkeley, and a research associate of the National Bureau of Economic Research.

the literature. It is divided into micro- and macroeconomic estimates at three levels: (1) the returns to the individual; (2) the returns to the nation, which may well be higher if there are externalities; and (3) the returns to the world, which may be higher still if there are foreign externalities. Of course, there is also the possibility of negative externalities, and Dowrick considers the costs of extra education which also play into the analysis.

In future work like this, I would like to see more emphasis placed on the externalities themselves. How much do we really know from quantifiable microeconomic evidence on the existence of large positive externalities? Since this is where the real case for policy intervention lies, I personally would feel more confident if I could cite a number of reliable studies that present strong evidence of positive externalities.

The reason I would find this is reassuring lies in the magnitudes of the returns to education cited. Almost all the returns to education and R&D are high—huge, in fact. It makes me feel that I've personally underinvested! More generally, the returns are so high that they strain plausibility. Lots of education is wasted, and much R&D might well be unproductive—is it all being taken into account? Let me put it another way. The returns are so large that the question of the paper's title is almost irrelevant, since the issue of levels versus growth rate effects is a sideshow if the returns are so high. So is concern for underinvestment in education, if the personal returns are as high as cited.

I also believe there is a lot of scope in this area for a comprehensive meta-analysis. This is the increasingly accepted way to conduct a quantitative survey. The author chooses a coefficient of interest that has been estimated in a number of papers—for instance, the value of a marginal year of education. Each paper contributes a single observation of this underlying variable, and the resulting vector of estimates is treated as the dependent variable. The characteristics of the studies are treated as the regressors. Meta-analysis like this might enable us to understand the sources of variation in the estimates of the returns to education, and enable us to handle them with more confidence.