

Jones, Charles (2002) *Introduction to Economic Growth* 2nd ed (W.W. Norton)

Chapter 3: Empirical Applications of Neoclassical Growth Models



his chapter considers several applications of the Solow model and its descendents, which we will group together under the rubric of “neoclassical growth models.” In the first section of this chapter, we develop one of the key descendents of the Solow model, an extension that incorporates human capital. Then, we examine the “fit” of the model: How well does the neoclassical growth model explain why some countries are rich and others are poor? In the second section of this chapter, we examine the model’s predictions concerning growth rates and discuss the presence or lack of “convergence” in the data. Finally, the third section of this chapter merges the discussion of the cross-country distribution of income levels with the convergence literature and examines the evolution of the world income distribution.

3.1 THE SOLOW MODEL WITH HUMAN CAPITAL

In an influential paper published in 1992, “A Contribution to the Empirics of Economic Growth,” Gregory Mankiw, David Romer, and David Weil evaluated the empirical implications of the Solow model and con-

cluded that it performed very well. They then noted that the “fit” of the model could be improved even more by extending the model to include human capital—that is, by recognizing that labor in different economies may possess different levels of education and different skills. Extending the Solow model to include human capital or skilled labor is relatively straightforward, as we shall see in this section.¹

Suppose that output, Y , in an economy is produced by combining physical capital, K , with skilled labor, H , according to a constant-returns, Cobb-Douglas production function

$$Y = K^\alpha (AH)^{1-\alpha}, \quad (3.1)$$

where A represents labor-augmenting technology that grows exogenously at rate g .

Individuals in this economy accumulate human capital by spending time learning new skills instead of working. Let u denote the fraction of an individual’s time spent learning skills, and let L denote the total amount of (raw) labor used in production in the economy.² We assume that unskilled labor learning skills for time u generates skilled labor H according to

$$H = e^{\psi u} L, \quad (3.2)$$

where ψ is a positive constant we will discuss in a moment. Notice that if $u = 0$, then $H = L$ —that is, all labor is unskilled. By increasing u , a unit of unskilled labor increases the *effective* units of skilled labor H . To see by how much, take logs and derivatives of equation (3.2) to see that

$$\frac{d \log H}{du} = \psi \implies \frac{dH}{du} = \psi H. \quad (3.3)$$

To interpret this equation, suppose that u increases by 1 unit (think of this as one additional year of schooling), and suppose $\psi = .10$. In this

¹The development here differs from that in Mankiw, Romer, and Weil (1992) in one important way. Mankiw, Romer, and Weil allow an economy to accumulate human capital in the same way that it accumulates physical capital: by foregoing consumption. Here, instead, we follow Lucas (1988) in assuming that individuals spend time accumulating skills, much like a student going to school. See Exercise 6 at the end of this chapter.

²Notice that if P denotes the total population of the economy, then the total amount of labor input in the economy is given by $L \equiv (1 - u)P$.

case, H rises by 10 percent. The fact that the effects are proportional is driven by the somewhat odd presence of the exponential e in the equation. This formulation is intended to match a large literature in labor economics that finds that an additional year of schooling increases the wages earned by an individual by something like 10 percent.³

Physical capital is accumulated by investing some output instead of consuming it, as in Chapter 2:

$$\dot{K} = s_K Y - dK, \quad (3.4)$$

where s_K is the investment rate for physical capital and d is the constant depreciation rate.

We solve this model using the same techniques employed in Chapter 2. First, we let lower-case letters denote variables divided by the stock of unskilled labor, L , and rewrite the production function in terms of output per worker as

$$y = k^\alpha (Ah)^{1-\alpha}. \quad (3.5)$$

Notice that $h = e^{\psi u}$. How do agents decide how much time to spend accumulating skills instead of working? Just as we assume that individuals save and invest a constant fraction of their income, we will assume that u is constant and given exogenously.⁴

The fact that h is constant means that the production function in equation (3.5) is very similar to that used in Chapter 2. In particular, along a balanced growth path, y and k will grow at the constant rate g , the rate of technological progress.

As in Chapter 2, the model is solved by considering “state variables” that are constant along a balanced growth path. There, recall that the state variables were terms such as y/A . Here, since h is constant, we can define the state variables by dividing by Ah . Denoting these state variables with a tilde, equation (3.5) implies that

$$\tilde{y} = \tilde{k}^\alpha, \quad (3.6)$$

which is the same as equation (2.11).

³Bils and Klenow (2000) apply this Mincerian formulation in the context of economic growth.

⁴We return to this issue in Chapter 7.

Following the reasoning from Chapter 2, the capital accumulation equation can be written in terms of the state variables as

$$\dot{\tilde{k}} = s_K \tilde{y} - (n + g + d)\tilde{k}. \quad (3.7)$$

Notice that in terms of state variables, this model is identical to the model we have already solved in Chapter 2. That is, equations (3.6) and (3.7) are identical to equations (2.11) and (2.12). This means that all of the results we discussed in Chapter 2 regarding the dynamics of the Solow model apply here. Adding human capital as we have done it does not change the basic flavor of the model.

The steady-state values of \tilde{k} and \tilde{y} are found by setting $\dot{\tilde{k}} = 0$, which yields

$$\frac{\tilde{k}}{\tilde{y}} = \frac{s_K}{n + g + d}.$$

Substituting this condition into the production function in equation (3.6), we find the steady-state value of the output-technology ratio \tilde{y} :

$$\tilde{y}^* = \left(\frac{s_K}{n + g + d} \right)^{\alpha/(1-\alpha)}.$$

Rewriting this in terms of output per worker, we get

$$y^*(t) = \left(\frac{s_K}{n + g + d} \right)^{\alpha/(1-\alpha)} hA(t), \quad (3.8)$$

where we have explicitly included t to remind us which variables are growing over time.

This last equation summarizes the explanation provided by the extended Solow model for why some countries are rich and others are poor. Countries are rich because they have high investment rates in physical capital, spend a large fraction of time accumulating skills ($h = e^{\psi u}$), have low population growth rates, and have high levels of technology. Furthermore, in the steady state, per capita output grows at the rate of technological progress, g , just as in the original Solow model.

How well does this model perform empirically in terms of explaining why some countries are richer than others? Because incomes are growing over time, it is useful to analyze the model in terms of *relative* incomes. If we define per capita income relative to the United States

to be

$$\hat{y}^* = \frac{y^*}{y_{US}^*},$$

then from equation (3.8), relative incomes are given by


$$\hat{y}^* = \left(\frac{\hat{s}_K}{\hat{x}} \right)^{\alpha/(1-\alpha)} \hat{h}\hat{A}, \quad (3.9)$$

where the “hat” ($\hat{\cdot}$) is used to denote a variable relative to its U.S. value, and $x \equiv n + g + d$. Notice, however, that unless countries are all growing at the same rate, even relative incomes will not be constant. That is, if the United Kingdom and the United States are growing at different rates, then y_{UK}/y_{US} will not be constant.

In order for relative incomes to be constant in the steady state, we need to make the assumption that g is the same in all countries—that is, the rate of technological progress in all countries is identical. On the surface, this seems very much at odds with one of our key stylized facts from Chapter 1: that growth rates vary substantially across countries. We will discuss technology in much greater detail in later chapters, but for now, notice that if g varies across countries, then the “income gap” between countries eventually becomes infinite. This may not seem plausible if growth is driven purely by technology. Technologies may flow across international borders through international trade, or in scientific journals and newspapers, or through the immigration of scientists and engineers. It may be more plausible to think that technology transfer will keep even the poorest countries from falling too far behind, and one way to interpret this statement is that the growth rates of technology, g , are the same across countries. We will formalize this argument in Chapter 6. In the meantime, notice that in no way are we requiring the *levels* of technology to be the same; in fact, differences in technology presumably help to explain why some countries are richer than others.

Still, we are left wondering why it is that countries have grown at such different rates over the last thirty years if they have the same underlying growth rate for technology. It may seem that the Solow model cannot answer this question, but in fact it provides a very good answer that will be discussed in the next section. First, however, we return to the basic question of how well the extended Solow model fits the data.

FIGURE 3.1 THE “FIT” OF THE NEOCLASSICAL GROWTH MODEL, 1997



Note: A log scale is used for each axis.

By obtaining estimates of the variables and parameters in equation (3.9), we can examine the “fit” of the neoclassical growth model: empirically, how well does it explain why some countries are rich and others are poor?

Figure 3.1 compares the actual levels of GDP per worker in 1997 to the levels predicted by equation (3.9). To use the equation, we assume a physical capital share of $\alpha = 1/3$. This choice fits well with the observation that the share of GDP paid to capital is about 1/3. We measure u as the average educational attainment of the labor force (in years) and assume that $\psi = .10$. Such a value implies that each year of schooling

increases a worker's wage by 10 percent, a number roughly consistent with international evidence on returns to schooling.⁵ In addition, we assume that $g + d = .075$ for all countries; we will discuss the assumption that g is the same in all countries in later chapters, and there is no good data on differences in d across countries. Finally, we assume that the technology level, A , is the same across countries. That is, we tie one hand behind our back to see how well the model performs without introducing technological differences. This assumption will be discussed shortly. The data used in this exercise are listed in Appendix C at the end of the book.

Without accounting for differences in technology, the neoclassical model still describes the distribution of per capita income across countries fairly well. Countries such as the United States and Norway are quite rich, as predicted by the model. Countries such as Uganda and Mozambique are decidedly poor. The main failure of the model — that it ignores differences in technology — can be seen by the departures from the 45-degree line in Figure 3.1: the model predicts that the poorest countries should be richer than they are.

How can we incorporate actual technology levels into the analysis? It is difficult to answer this question in a satisfactory manner, but there is a convenient “cheat” that is available. We can use the production function itself to solve for the level of A consistent with each country's output and capital. This is a cheat in that we are simply calculating A to make the model fit the data. However, it is an informative cheat. One can examine the A s that are needed to fit the data to see if they are plausible.


Solving the production function in Equation (3.5) for A yields

$$A = \left(\frac{y}{k}\right)^{\alpha/1-\alpha} \frac{y}{h}$$

With data on GDP per worker, capital per worker, and educational attainment for each country, we can use this equation to estimate actual levels of A . These estimates are reported in Figure 3.2.

⁵See Jones (1996) for additional details. Notice that measuring u as years of schooling means that it is no longer between zero and one. This problem can be addressed by dividing years of schooling by potential lifespan, which simply changes the value of ψ proportionally and is therefore ignored.

FIGURE 3.2 PRODUCTIVITY LEVELS, 1997



Note: A log scale is used for each axis, and U.S. values are normalized to 1.

From this figure, one discovers several important things. First, the levels of A calculated from the production function are strongly correlated with the levels of output per worker across countries. Rich countries generally have high levels of A , and poor countries generally have low levels. Countries that are rich not only have high levels of physical and human capital, but they also manage to use these inputs very productively.

Second, although levels of A are highly correlated with levels of income, the correlation is far from perfect. Countries such as Singapore, Italy, and Bangladesh have much higher levels of A than would be expected from their GDP per worker, and perhaps have levels that are too high to be plausible. Indeed, several countries have levels of A higher than that in the United States. This observation leads to an important

remark. Estimates of A computed this way are like the residuals from growth accounting: they incorporate *any* differences in production not factored in through the inputs. For example, we have not controlled for differences in the quality of educational systems, the importance of experience at work and on-the-job training, or the general health of the labor force. These differences will therefore be included in A . In this sense, it is more appropriate to refer to these estimates as total factor productivity levels rather than technology levels.

Finally, the differences in total factor productivity across countries are large. The poorest countries of the world have levels of A that are only 10 to 15 percent of those in the richest countries.

With this observation, we can return to equation (3.9) to make one last remark. The richest countries of the world have an output per worker that is roughly 32 times that of the poorest countries of the world. This difference can be broken down into differences associated with investment rates in physical capital, investment rates in human capital, and differences in productivity. For this purpose, it is helpful to refer to the data in Appendix C. The richest countries of the world have investment rates that are around 25 percent, while the poorest countries of the world have investment rates around 5 percent. As a rough approximation, then, s/x varies by about a factor of 5 across countries. According to equation (3.9), it is the square root of this factor (since $\alpha/1 - \alpha = 1/2$) that contributes to output per worker, so that differences in physical capital account for just over a factor of 2 of the differences in output per worker between the rich and poor countries.

Similarly, workers in rich countries have about 10 or 11 years of education on average, whereas workers in poor countries have less than 3 years. Assuming a return to schooling of 10 percent, this suggests that $\hat{h} \approx e^{-10(11-3)} \approx e^{-8} \approx 2.2$. That is, differences in educational attainment also contribute a factor of just over 2 to differences in output per worker between the rich and poor countries.

What accounts for the remainder? By construction, differences in total factor productivity contribute the remaining factor of 8 to the differences in output per worker between the rich and poor countries.⁶

⁶A more extensive analysis of productivity levels can be found in Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999).

Productivity differences across countries are large, and a satisfactory theory of growth and development needs to explain these differences.

In summary, the Solow framework is extremely successful in helping us to understand the wide variation in the wealth of nations. Countries that invest a large fraction of their resources in physical capital and in the accumulation of skills are rich. Countries that use these inputs productively are rich. The countries that fail in one or more of these dimensions suffer a corresponding reduction in income. Of course, one thing the Solow model does not help us understand is *why* some countries invest more than others, and *why* some countries attain higher levels of technology or productivity. Addressing these questions is the subject of Chapter 7. As a preview, the answers are tied intimately to government policies and institutions.

3.2 CONVERGENCE AND EXPLAINING DIFFERENCES IN GROWTH RATES

We have discussed in detail the ability of the neoclassical model to explain differences in income levels across economies, but how well does it perform at explaining differences in growth rates? An early hypothesis proposed by economic historians such as Aleksander Gerschenkron (1952) and Moses Abramovitz (1986) was that, at least under certain circumstances, “backward” countries would tend to grow faster than rich countries, in order to close the gap between the two groups. This catch-up phenomenon is referred to as *convergence*. For obvious reasons, questions about convergence have been at the heart of much empirical work on growth. We documented in Chapter 1 the enormous differences in levels of income per person around the world: the typical person in the United States earns in less than ten days the annual income of the typical person in Ethiopia. The question of convergence asks whether these enormous differences are getting smaller over time.

An important cause of convergence might be technology transfer, but the neoclassical growth model provides another explanation for convergence that we will explore in this section. First, however, let’s examine the empirical evidence on convergence.

William Baumol (1986), alert to the analysis provided by economic historians, was one of the first economists to provide statistical evidence documenting convergence among some countries and the absence of convergence among others. The first piece of evidence presented by Baumol is displayed in Figure 3.3, which plots per capita GDP (on a log scale) for several industrialized economies from 1870 to 1994. The narrowing of the gaps between countries is evident in this figure. Interestingly, the world “leader” in terms of per capita GDP in 1870 was Australia (not shown). The United Kingdom had the second-highest per capita GDP and was recognized as the industrial center of the Western world. Around the turn of the century, the United States surpassed Australia and the United Kingdom and has remained the “leader” ever since.

FIGURE 3.3 PER CAPITA GDP, 1870-1994




FIGURE 3.4 GROWTH RATE VERSUS INITIAL PER CAPITA GDP, 1885-1994





Figure 3.4 reveals the ability of the convergence hypothesis to explain why some countries grew fast and others grew slowly over the course of the last century. The graph plots a country's initial per capita GDP (in 1885) against the country's growth rate from 1885 to 1994. The figure reveals a strong negative relationship between the two variables: countries such as Australia and the United Kingdom, which were relatively rich in 1885, grew most slowly, while countries like Japan that were relatively poor grew most rapidly. The simple convergence hypothesis seems to do a good job of explaining differences in growth rates, at least among this sample of industrialized economies.⁷

Figures 3.5 and 3.6 plot growth rates versus initial GDP per worker for the countries that are members of the Organization for Economic

⁷J. Bradford DeLong (1988) provides an important criticism of this result. See Exercise 5 at the end of this chapter.


FIGURE 3.5 CONVERGENCE IN THE OECD, 1960-97



Cooperation and Development (OECD) and for the world for the period 1960–97. Figure 3.5 shows that the convergence hypothesis works extremely well for explaining growth rates across the OECD for the period examined. But before we declare the hypothesis a success, note that Figure 3.6 shows that the convergence hypothesis fails to explain differences in growth rates across the world as a whole. Baumol also reported this finding: across large samples of countries, it does not appear that poor countries grow faster than rich countries. The poor countries are not “closing the gap” that exists in per capita incomes. (Recall that Table 1.1 in Chapter 1 supports this finding.)

Why, then, do we see convergence among some sets of countries but a lack of convergence among the countries of the world as a whole? The neoclassical growth model suggests an important explanation for these findings.

FIGURE 3.6 THE LACK OF CONVERGENCE FOR THE WORLD, 1960-97




Consider the key differential equation of the neoclassical growth model, given in equation (3.7). This equation can be rewritten as

$$\frac{\dot{\bar{k}}}{\bar{k}} = s_K \frac{\bar{y}}{\bar{k}} - (n + g + d). \quad (3.10)$$

Remember that \bar{y} is equal to \bar{k}^α . Therefore, the average product of capital \bar{y}/\bar{k} is equal to $\bar{k}^{\alpha-1}$. In particular, it declines as \bar{k} rises, because of the diminishing returns to capital accumulation in the neoclassical model.

As in Chapter 2, we can analyze this equation in a simple diagram, shown in Figure 3.7. The two curves in the figure plot the two terms on the right-hand side of equation (3.10). Therefore, the difference between the curves is the growth rate of \bar{k} . Notice that the growth rate of \bar{y} is simply proportional to this difference. Furthermore, because the growth

FIGURE 3.7 TRANSITION DYNAMICS IN THE NEOCLASSICAL MODEL



rate of technology is constant, any changes in the growth rates of \tilde{k} and \tilde{y} must be due to changes in the growth rates of capital per worker, k , and output per worker, y .

Suppose the economy of InitiallyBehind starts with the capital-technology ratio \tilde{k}_{IB} shown on Figure 3.7, while the neighboring economy of InitiallyAhead starts with the higher capital-technology ratio indicated by \tilde{k}_{IA} . If these two economies have the same levels of technology, the same rates of investment, and the same rates of population growth, then InitiallyBehind will temporarily grow faster than InitiallyAhead. The output-per-worker gap between the two countries will narrow over time as both economies approach the same steady state. An important prediction of the neoclassical model is this: *Among countries that have the same steady state, the convergence hypothesis should hold: poor countries should grow faster on average than rich countries.*

For the industrialized countries, the assumption that their economies have similar technology levels, investment rates, and population

growth rates may not be a bad one. The neoclassical model, then, would predict the convergence that we saw in Figures 3.4 and 3.5. This same reasoning suggests a compelling explanation for the *lack* of convergence across the world as a whole: all countries do not have the same steady states. In fact, as we saw in Figure 3.2, the differences in income levels around the world largely reflect differences in steady states. Because all countries do not have the same investment rates, population growth rates, or technology levels, they are not generally expected to grow toward the same steady-state target.

Another important prediction of the neoclassical model is related to growth rates. This prediction, which can be found in many growth models, is important enough that we will give it a name, the “principle of transition dynamics”:


The further an economy is “below” its steady state, the faster the economy should grow. The further an economy is “above” its steady state, the slower the economy should grow.⁸

This principle is clearly illustrated by the analysis of equation (3.10) provided in Figure 3.7. Although it is a key feature of the neoclassical model, the principle of transition dynamics applies much more broadly. In Chapters 5 and 6, for example, we will see that it is also a feature of the models of new growth theory that endogenize technological progress.

Mankiw, Romer, and Weil (1992) and Barro and Sala-i-Martin (1992) show that this prediction of the neoclassical model can explain differences in growth rates across the countries of the world. Figure 3.8 illustrates this point by plotting the growth rate of GDP per worker from 1960 to 1997 against the deviation of GDP per worker (relative to that of the U.S.) from its steady-state value. This steady state is computed according to equation (3.9) using the data in Appendix C and a total factor productivity level from 1970. (You will be asked to undertake a similar calculation in Exercise 1 at the end of the chapter.) Comparing Figure 3.6 and Figure 3.8, one sees that although poorer countries do not necessarily grow faster, countries that are “poor” relative to their own steady states do tend to grow more rapidly. In 1960, good

⁸In simple models, including most of those presented in this book, this principle works well. In more complicated models with more state variables, however, it must be modified.

FIGURE 3.3 “CONDITIONAL” CONVERGENCE FOR THE WORLD, 1960–97



Note: The variable on the x-axis is \hat{y}_{60}/\hat{y}^* . Estimates of A for 1970 are used to compute the steady state.

examples of these countries were Korea, Japan, Singapore, and Hong Kong — economies that grew rapidly over the next forty years, just as the neoclassical model would predict.⁹

⁹Mankiw, Romer, and Weil (1992) and Barro and Sala-i-Martin (1992) have called this phenomenon “conditional convergence,” because it reflects the convergence of countries after we control for (“condition on”) differences in steady states. It is important to keep in mind what this “conditional convergence” result means. It is simply a confirmation of a result predicted by the neoclassical growth model: that countries with similar steady states will exhibit convergence. It does not mean that all countries in the world are converging to the same steady state, only that they are converging to their own steady states according to a common theoretical model.

This analysis of convergence has been extended by a number of authors to different sets of economies. For example, Barro and Sala-i-Martin (1991, 1992) show that the U.S. states, regions of France, and prefectures in Japan all exhibit “unconditional” convergence similar to what we’ve observed in the OECD. This matches the prediction of the Solow model if regions within a country are similar in terms of investment and population growth, as seems reasonable.

How does the neoclassical model account for the wide differences in growth rates across countries documented in Chapter 1? The principle of transition dynamics provides the answer: countries that have not reached their steady states are not expected to grow at the same rate. Those “below” their steady states will grow rapidly, and those “above” their steady states will grow slowly.


As we saw in Chapter 2, there are many reasons why countries may not be in steady state. An increase in the investment rate, a change in the population growth rate, or an event like World War II that destroys much of a country’s capital stock will generate a gap between current income and steady-state income. This gap will change growth rates until the economy returns to its steady-state path. Other “shocks” can also cause temporary differences in growth rates. For example, large changes in oil prices will have important effects on the economic performance of oil-exporting countries. Mismanagement of the macroeconomy can similarly generate temporary changes in growth performance. The hyperinflations in many Latin American countries during the 1980s are a good example of this. Working in the other direction, policy reforms that shift the steady-state path of an economy upward can generate increases in growth rates along a transition path. Increases in the investment rate, skill accumulation, or the level of technology will have this effect.¹⁰

3.3 THE EVOLUTION OF THE INCOME DISTRIBUTION

Convergence, the closing of the gap between rich and poor economies, is just one possible outcome among many that could be occurring. Alternatively, perhaps the poorest countries are falling behind while countries

¹⁰Barro (1991) and Easterly, Kremer, et al. (1993) provide empirical analyses of why countries have exhibited different growth rates since 1960.

FIGURE 3.9
INCOME RATIOS, 5TH-RICHEST COUNTRY TO
5TH-POOREST COUNTRY, 1960-97



with “intermediate” incomes are converging toward the rich. Or perhaps countries are not getting any closer together at all but are instead fanning out, with the rich countries getting richer and the poor countries getting poorer. More generally, these questions are really about the evolution of the distribution of per capita incomes around the world.¹¹

Figure 3.9 illustrates a key fact about the evolution of the income distribution: for the world as a whole, the enormous gaps in income across countries have generally not narrowed over time. This figure

plots the ratio of GDP per worker in the 5th-richest country to GDP per worker in the 5th-poorest country. In 1960, GDP per worker of the fifth-richest country was more than 25 times that of the fifth-poorest country. By 1990, the ratio had risen slightly, to around 30. The 1990s witnessed an even sharper increase, to more than 35 by the end of the sample.

The widening of the world income distribution is a fact that almost certainly characterizes the world economy over its entire history. Incomes cannot get much lower than about \$250: below this level widespread starvation and death set in. This number provides a lower bound on incomes at any date in the past, and this lower bound comes close to being attained by the poorest countries in the world even today. On the other hand, the incomes of the richest countries have been growing over time. This suggests that the ratio of the incomes in the richest to those in the poorest countries has also been rising. Lant Pritchett (1997), in a paper titled “Divergence, Big Time,” calculates that the ratio of per capita GDP between the richest and poorest countries in the world was only 8.7 in 1870 but rose to 45.2 by 1990. Before 1870, the ratio was presumably even lower.


Whether this widening will continue in the future is an open question. One possible explanation for the increase is that countries climb onto the modern economic growth “escalator” at different points in time. As long as there are some countries that have yet to get on, the world income distribution widens. Once all countries get on, however, this widening may reverse.¹²

While Figure 3.9 shows that the “width” of the income distribution has increased, Figure 3.10 examines changes at each point in the income distribution. According to the figure, 50 percent of the countries had relative incomes that were less than 15 percent of U.S. GDP per worker in 1960; 80 percent of the countries had relative incomes less than 40 percent of U.S. GDP per worker. By 1997, these numbers had improved, particularly at the upper end: the 50th percentile was slightly less than 20 percent of U.S. GDP per worker while the 80th percentile was more than 60 percent. In contrast, the poorest economies — those below the 30th percentile, for example — actually had relative incomes in 1997 lower than in 1960. In this sense, one might say there was some catch-up or convergence at the middle and top of the income distribu-

¹¹Jones (1997) provides an overview of the literature on the world income distribution. Quah (1993, 1996) discusses this topic in more detail.

¹²Robert E. Lucas, Jr. (2000), analyzes a model like this in a very readable manner.

FIGURE 3.10 THE EVOLUTION OF THE WORLD INCOME DISTRIBUTION



Note: A point (x, y) in the figure indicates that x percent of countries had relative GDP per worker less than or equal to y . One hundred ten countries are represented.

tion from 1960 to 1997, but divergence at the bottom end.¹³ Danny Quah (1996) suggests that this tendency for the middle-income countries to become relatively richer while the poorest countries become relatively (but not necessarily absolutely) poorer will result in an income distribution with “twin peaks”—i.e., a mass of countries at both ends of the income distribution.

¹³It is interesting to compare this figure to the results in Chapter 1. An important difference is that the unit of observation here is the *country*; the unit of observation for the distributions computed in Chapter 1 was the *individual*.

EXERCISES

1. *Where are these economies headed?* Consider the following data:

	\hat{y}_{97}	s_K	u	n	\hat{A}_{90}
U.S.A.	1.000	0.204	11.9	0.010	1.000
Canada	0.864	0.246	11.4	0.012	0.972
Argentina	0.453	0.144	8.5	0.014	0.517
Thailand	0.233	0.213	6.1	0.015	0.468
Cameroon	0.048	0.102	3.4	0.028	0.234

Assume that $g + d = .075$, $\alpha = 1/3$, and $\psi = .10$ for all countries. Using equation (3.9), estimate the steady-state incomes of these economies, relative to the United States. Consider two extreme cases: (a) the 1990 TFP ratios are maintained, and (b) the TFP levels converge completely. For each case, which economy will grow fastest in the next decade and which slowest? Why?

2. *Policy reforms and growth.* Suppose an economy, starting from an initial steady state, undertakes new policy reforms that raise its steady-state level of output per worker. For each of the following cases, calculate the proportion by which steady-state output per worker increases and, using the slope of the relationship shown in Figure 3.8, make a guess as to the amount by which the growth rate of GDP per worker will be higher during the next forty years. Assume $\alpha = 1/3$ and $\psi = .10$. (a) The level of total factor productivity, A is permanently doubled. (b) The investment rate, s_K , is permanently doubled. (c) The average educational attainment of the labor force, u , is permanently increased by 5 years.
3. *What are state variables?* The basic idea of solving dynamic models that contain a differential equation is to first write the model so that along a balanced growth path, some state variable is constant. In Chapter 2, we used y/A and k/A as state variables. In this chapter, we used y/Ah and k/Ah . Recall, however, that h is a constant. This reasoning suggests that one should be able to solve the model using y/A and k/A as the state variables. Do this. That is, solve the growth model in equations (3.1) to (3.4) to get the solution in equation (3.8) using y/A and k/A as state variables.

4. *Galton's fallacy* (based on Quah 1993). During the late 1800s, Sir Francis Galton, a famous statistician in England, studied the distribution of heights in the British population and how the distribution was evolving over time. In particular, Galton noticed that the sons of tall fathers tended to be shorter than their fathers, and vice versa. Galton worried that this implied some kind of regression toward "mediocrity."

Suppose that we have a population of 10 mothers who have 10 daughters. Suppose that their heights are determined as follows. Place 10 sheets of paper in a hat labeled with heights of 5'1", 5'2", 5'3", . . . 5'10". Draw a number from the hat and let that be the height for a mother. Without replacing the sheet just drawn, continue. Now suppose that the heights of the daughters are determined in the same way, starting with the hat full again and drawing new heights. Make a graph of the change in height between daughter and mother against the height of the mother. Will tall mothers tend to have shorter daughters, and vice versa?

Let the heights correspond to income levels, and consider observing income levels at two points in time, say 1960 and 1990. What does Galton's fallacy imply about a plot of growth rates against initial income? Does this mean the figures in this chapter are useless?¹⁴

5. *Reconsidering the Baumol results*. J. Bradford DeLong (1988), in a comment on Baumol's convergence result for the industrialized countries over the last century, pointed out that the result could be driven by the procedure through which the countries were selected. In particular, DeLong noted two things. First, only countries that were rich at the end of the sample (i.e., in the 1980s) were included. Second, several countries not included, such as Argentina, were richer than Japan in 1870. Use these points to criticize and discuss the Baumol results. Do these criticisms apply to the results for the OECD? For the world?
6. *The Mankiw-Romer-Weil (1992) model*. As mentioned in this chapter, the extended Solow model that we have considered differs slightly from that in Mankiw, Romer, and Weil (1992). This problem asks you to solve their model. The key difference is the treatment of hu-

man capital. Mankiw, Romer, and Weil assume that human capital is accumulated just like physical capital, so that it is measured in units of output instead of years of time.

Assume production is given by $Y = K^\alpha H^\beta (AL)^{1-\alpha-\beta}$, where α and β are constants between zero and one whose sum is also between zero and one. Human capital is accumulated just like physical capital:

$$\dot{H} = s_H Y - dH,$$

where s_H is the constant share of output invested in human capital. Assume that physical capital is accumulated as in equation (3.4), that the labor force grows at rate n , and that technological progress occurs at rate g . Solve the model for the path of output per worker $y \equiv Y/L$ along the balanced growth path as a function of s_K , s_H , n , g , d , α , and β . Discuss how the solution differs from that in equation (3.8). (Hint: Define state variables such as y/A , h/A , and k/A .)

¹⁴See Quah (1993) and Friedman (1992).