



Economic Growth II: Technology, Empirics, and Policy

Is there some action a government of India could take that would lead the Indian economy to grow like Indonesia's or Egypt's? If so, what, exactly? If not, what is it about the "nature of India" that makes it so? The consequences for human welfare involved in questions like these are simply staggering: Once one starts to think about them, it is hard to think about anything else.

—Robert E. Lucas, Jr., 1988

This chapter continues our analysis of the forces governing long-run economic growth. With the basic version of the Solow growth model as our starting point, we take on four new tasks.

Our first task is to make the Solow model more general and realistic. In Chapter 3 we saw that capital, labor, and technology are the key determinants of a nation's production of goods and services. In Chapter 7 we developed the Solow model to show how changes in capital (through saving and investment) and changes in the labor force (through population growth) affect the economy's output. We are now ready to add the third source of growth—changes in technology—to the mix. The Solow model does not explain technological progress but, instead, takes it as exogenously given and shows how it interacts with other variables in the process of economic growth.

Our second task is to move from theory to empirics. That is, we consider how well the Solow model fits the facts. Over the past two decades, a large literature has examined the predictions of the Solow model and other models of economic growth. It turns out that the glass is both half full and half empty. The Solow model can shed much light on international growth experiences, but it is far from the last word on the subject.

Our third task is to examine how a nation's public policies can influence the level and growth of its citizens' standard of living. In particular, we address five questions: Should our society save more or less? How can policy influence the rate of saving? Are there some types of investment that policy should especially encourage? What institutions ensure that the economy's resources are put to their best use? How can policy increase the rate of technological progress? The

Solow growth model provides the theoretical framework within which we consider these policy issues.

Our fourth and final task is to consider what the Solow model leaves out. As we have discussed previously, models help us understand the world by simplifying it. After completing an analysis of a model, therefore, it is important to consider whether we have oversimplified matters. In the last section, we examine a new set of theories, called *endogenous growth theories*, which help to explain the technological progress that the Solow model takes as exogenous.

8-1 Technological Progress in the Solow Model

So far, our presentation of the Solow model has assumed an unchanging relationship between the inputs of capital and labor and the output of goods and services. Yet the model can be modified to include exogenous technological progress, which over time expands society's production capabilities.

The Efficiency of Labor

To incorporate technological progress, we must return to the production function that relates total capital K and total labor L to total output Y . Thus far, the production function has been

$$Y = F(K, L).$$

We now write the production function as

$$Y = F(K, L \times E),$$

where E is a new (and somewhat abstract) variable called the **efficiency of labor**. The efficiency of labor is meant to reflect society's knowledge about production methods: as the available technology improves, the efficiency of labor rises, and each hour of work contributes more to the production of goods and services. For instance, the efficiency of labor rose when assembly-line production transformed manufacturing in the early twentieth century, and it rose again when computerization was introduced in the late twentieth century. The efficiency of labor also rises when there are improvements in the health, education, or skills of the labor force.

The term $L \times E$ can be interpreted as measuring the *effective number of workers*. It takes into account the number of actual workers L and the efficiency of each worker E . In other words, L measures the number of workers in the labor force, whereas $L \times E$ measures both the workers and the technology with which the typical worker comes equipped. This new production function states that total output Y depends on the inputs of capital K and effective workers $L \times E$.

The essence of this approach to modeling technological progress is that increases in the efficiency of labor E are analogous to increases in the labor

force L . Suppose, for example, that an advance in production methods makes the efficiency of labor E double between 1980 and 2010. This means that a single worker in 2010 is, *in effect*, as productive as two workers were in 1980. That is, even if the actual number of workers (L) stays the same from 1980 to 2010, the effective number of workers ($L \times E$) doubles, and the economy benefits from the increased production of goods and services.

The simplest assumption about technological progress is that it causes the efficiency of labor E to grow at some constant rate g . For example, if $g = 0.02$, then each unit of labor becomes 2 percent more efficient each year: output increases as if the labor force had increased by 2 percent more than it really did. This form of technological progress is called *labor augmenting*, and g is called the rate of **labor-augmenting technological progress**. Because the labor force L is growing at rate n , and the efficiency of each unit of labor E is growing at rate g , the effective number of workers $L \times E$ is growing at rate $n + g$.

The Steady State With Technological Progress

Because technological progress is modeled here as labor augmenting, it fits into the model in much the same way as population growth. Technological progress does not cause the actual number of workers to increase, but because each worker in effect comes with more units of labor over time, technological progress causes the effective number of workers to increase. Thus, the analytic tools we used in Chapter 7 to study the Solow model with population growth are easily adapted to studying the Solow model with labor-augmenting technological progress.

We begin by reconsidering our notation. Previously, when there was no technological progress, we analyzed the economy in terms of quantities per worker; now we can generalize that approach by analyzing the economy in terms of quantities per effective worker. We now let $k = K/(L \times E)$ stand for capital per effective worker and $y = Y/(L \times E)$ stand for output per effective worker. With these definitions, we can again write $y = f(k)$.


Our analysis of the economy proceeds just as it did when we examined population growth. The equation showing the evolution of k over time becomes

$$\Delta k = sf(k) - (\delta + n + g)k.$$

As before, the change in the capital stock Δk equals investment $sf(k)$ minus break-even investment $(\delta + n + g)k$. Now, however, because $k = K/(L \times E)$, break-even investment includes three terms: to keep k constant, δk is needed to replace depreciating capital, nk is needed to provide capital for new workers, and gk is needed to provide capital for the new “effective workers” created by technological progress.¹

¹ *Mathematical note:* This model with technological progress is a strict generalization of the model analyzed in Chapter 7. In particular, if the efficiency of labor is constant at $E = 1$, then $g = 0$, and the definitions of k and y reduce to our previous definitions. In this case, the more general model considered here simplifies precisely to the Chapter 7 version of the Solow model.

FIGURE 8-1

Investment,
break-even
investment**Technological Progress and the Solow Growth Model**

Labor-augmenting technological progress at rate g enters our analysis of the Solow growth model in much the same way as did population growth at rate n . Now that k is defined as the amount of capital per effective worker, increases in the effective number of workers because of technological progress tend to decrease k . In the steady state, investment $sf(k)$ exactly offsets the reductions in k attributable to depreciation, population growth, and technological progress.

As shown in Figure 8-1, the inclusion of technological progress does not substantially alter our analysis of the steady state. There is one level of k , denoted k^* , at which capital per effective worker and output per effective worker are constant. As before, this steady state represents the long-run equilibrium of the economy.

The Effects of Technological Progress

Table 8-1 shows how four key variables behave in the steady state with technological progress. As we have just seen, capital per effective worker k is constant in the steady state. Because $y = f(k)$, output per effective worker is also constant. It is these quantities per effective worker that are steady in the steady state.

From this information, we can also infer what is happening to variables that are not expressed in units per effective worker. For instance, consider output per actual worker $Y/L = y \times E$. Because y is constant in the steady state and E is growing at rate g , output per worker must also be growing at rate g in the steady state. Similarly, the economy's total output is $Y = y \times (E \times L)$. Because y is constant in the steady state, E is growing at rate g , and L is growing at rate n , total output grows at rate $n + g$ in the steady state.

With the addition of technological progress, our model can finally explain the sustained increases in standards of living that we observe. That is, we have shown that technological progress can lead to sustained growth in output per worker. By contrast, a high rate of saving leads to a high rate of growth only until the steady state is reached. Once the economy is in steady state, the rate of growth of output per worker depends only on the rate of technological progress. *According to the Solow model, only technological progress can explain sustained growth and persistently rising living standards.*

TABLE 8-1

Steady-State Growth Rates in the Solow Model
With Technological Progress

Variable	Symbol	Steady-State Growth Rate
Capital per effective worker	$k = K/(E \times L)$	0
Output per effective worker	$y = Y/(E \times L) = f(k)$	0
Output per worker	$Y/L = y \times E$	g
Total output	$Y = y \times (E \times L)$	$n + g$

The introduction of technological progress also modifies the criterion for the Golden Rule. The Golden Rule level of capital is now defined as the steady state that maximizes consumption per effective worker. Following the same arguments that we have used before, we can show that steady-state consumption per effective worker is

$$c^* = f(k^*) - (\delta + n + g)k^*.$$

Steady-state consumption is maximized if

$$MPK = \delta + n + g,$$

or

$$MPK - \delta = n + g.$$

That is, at the Golden Rule level of capital, the net marginal product of capital, $MPK - \delta$, equals the rate of growth of total output, $n + g$. Because actual economies experience both population growth and technological progress, we must use this criterion to evaluate whether they have more or less capital than they would at the Golden Rule steady state.

8-2 From Growth Theory to Growth Empirics

So far in this chapter we have introduced exogenous technological progress into the Solow model to explain sustained growth in standards of living. Let's now discuss what happens when this theory is forced to confront the facts.

Balanced Growth

According to the Solow model, technological progress causes the values of many variables to rise together in the steady state. This property, called *balanced growth*, does a good job of describing the long-run data for the U.S. economy.

Consider first output per worker Y/L and the capital stock per worker K/L . According to the Solow model, in the steady state, both of these variables grow at g , the rate of technological progress. U.S. data for the past half century show that output per worker and the capital stock per worker have in fact grown at approximately the same rate—about 2 percent per year. To put it another way, the capital–output ratio has remained approximately constant over time.

Technological progress also affects factor prices. Problem 3(d) at the end of the chapter asks you to show that, in the steady state, the real wage grows at the rate of technological progress. The real rental price of capital, however, is constant over time. Again, these predictions hold true for the United States. Over the past 50 years, the real wage has increased about 2 percent per year; it has increased at about the same rate as real GDP per worker. Yet the real rental price of capital (measured as real capital income divided by the capital stock) has remained about the same.

The Solow model’s prediction about factor prices—and the success of this prediction—is especially noteworthy when contrasted with Karl Marx’s theory of the development of capitalist economies. Marx predicted that the return to capital would decline over time and that this would lead to economic and political crisis. Economic history has not supported Marx’s prediction, which partly explains why we now study Solow’s theory of growth rather than Marx’s.

Convergence

If you travel around the world, you will see tremendous variation in living standards. The world’s poor countries have average levels of income per person that are less than one-tenth the average levels in the world’s rich countries. These differences in income are reflected in almost every measure of the quality of life—from the number of televisions and telephones per household to the infant mortality rate and life expectancy.

Much research has been devoted to the question of whether economies converge over time to one another. In particular, do economies that start off poor subsequently grow faster than economies that start off rich? If they do, then the world’s poor economies will tend to catch up with the world’s rich economies. This property of catch-up is called *convergence*. If convergence does not occur, then countries that start off behind are likely to remain poor.

The Solow model makes clear predictions about when convergence should occur. According to the model, whether two economies will converge depends on why they differ in the first place. On the one hand, suppose two economies happen by historical accident to start off with different capital stocks, but they have the same steady state, as determined by their saving rates, population growth rates, and efficiency of labor. In this case, we should expect the two economies to converge; the poorer economy with the smaller capital stock will naturally grow more quickly to reach the steady state. (In a case study in Chapter 7, we applied this logic to explain rapid growth in Germany and Japan after World War II.) On the other hand, if two economies have different steady states, perhaps because the economies have different rates of saving, then we should not expect convergence. Instead, each economy will approach its own steady state.

Experience is consistent with this analysis. In samples of economies with similar cultures and policies, studies find that economies converge to one another at a rate of about 2 percent per year. That is, the gap between rich and poor economies closes by about 2 percent each year. An example is the economies of individual American states. For historical reasons, such as the Civil War of the 1860s, income levels varied greatly among states at the end of the nineteenth century. Yet these differences have slowly disappeared over time.

In international data, a more complex picture emerges. When researchers examine only data on income per person, they find little evidence of convergence: countries that start off poor do not grow faster on average than countries that start off rich. This finding suggests that different countries have different steady states. If statistical techniques are used to control for some of the determinants of the steady state, such as saving rates, population growth rates, and accumulation of human capital (education), then once again the data show convergence at a rate of about 2 percent per year. In other words, the economies of the world exhibit *conditional convergence*: they appear to be converging to their own steady states, which in turn are determined by such variables as saving, population growth, and human capital.²

Factor Accumulation Versus Production Efficiency

As a matter of accounting, international differences in income per person can be attributed to either (1) differences in the factors of production, such as the quantities of physical and human capital, or (2) differences in the efficiency with which economies use their factors of production. That is, a worker in a poor country may be poor because he lacks tools and skills or because the tools and skills he has are not being put to their best use. To describe this issue in terms of the Solow model, the question is whether the large gap between rich and poor is explained by differences in capital accumulation (including human capital) or differences in the production function.

Much research has attempted to estimate the relative importance of these two sources of income disparities. The exact answer varies from study to study, but both factor accumulation and production efficiency appear important. Moreover, a common finding is that they are positively correlated: nations with high levels of physical and human capital also tend to use those factors efficiently.³

There are several ways to interpret this positive correlation. One hypothesis is that an efficient economy may encourage capital accumulation. For example, a

² Robert Barro and Xavier Sala-i-Martin, “Convergence Across States and Regions,” *Brookings Papers on Economic Activity* 1 (1991): 107–182; and N. Gregory Mankiw, David Romer, and David N. Weil, “A Contribution to the Empirics of Economic Growth,” *Quarterly Journal of Economics* (May 1992): 407–437.

³ Robert E. Hall and Charles I. Jones, “Why Do Some Countries Produce So Much More Output per Worker Than Others?” *Quarterly Journal of Economics* 114 (February 1999): 83–116; and Peter J. Klenow and Andres Rodriguez-Clare, “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?” *NBER Macroeconomics Annual* (1997): 73–103.

person in a well-functioning economy may have greater resources and incentive to stay in school and accumulate human capital. Another hypothesis is that capital accumulation may induce greater efficiency. If there are positive externalities to physical and human capital, then countries that save and invest more will appear to have better production functions (unless the research study accounts for these externalities, which is hard to do). Thus, greater production efficiency may cause greater factor accumulation, or the other way around.

A final hypothesis is that both factor accumulation and production efficiency are driven by a common third variable. Perhaps the common third variable is the quality of the nation's institutions, including the government's policymaking process. As one economist put it, when governments screw up, they screw up big time. Bad policies, such as high inflation, excessive budget deficits, widespread market interference, and rampant corruption, often go hand in hand. We should not be surprised that economies exhibiting these maladies both accumulate less capital and fail to use the capital they have as efficiently as they might.

CASE STUDY

Is Free Trade Good for Economic Growth?

At least since Adam Smith, economists have advocated free trade as a policy that promotes national prosperity. Here is how Smith put the argument in his 1776 classic, *The Wealth of Nations*:

It is a maxim of every prudent master of a family, never to attempt to make at home what it will cost him more to make than to buy. The tailor does not attempt to make his own shoes, but buys them of the shoemaker. The shoemaker does not attempt to make his own clothes but employs a tailor. . . .

What is prudence in the conduct of every private family can scarce be folly in that of a great kingdom. If a foreign country can supply us with a commodity cheaper than we ourselves can make it, better buy it of them with some part of the produce of our own industry employed in a way in which we have some advantage.

Today, economists make the case with greater rigor, relying on David Ricardo's theory of comparative advantage as well as more modern theories of international trade. According to these theories, a nation open to trade can achieve greater production efficiency and a higher standard of living by specializing in those goods for which it has a comparative advantage.

A skeptic might point out that this is just a theory. What about the evidence? Do nations that permit free trade in fact enjoy greater prosperity? A large body of literature addresses precisely this question.

One approach is to look at international data to see if countries that are open to trade typically enjoy greater prosperity. The evidence shows that they do. Economists Andrew Warner and Jeffrey Sachs studied this question for the period from 1970 to 1989. They report that among developed nations, the open economies grew at 2.3 percent per year, while the closed economies grew at 0.7 percent per year. Among developing nations, the open economies grew at 4.5

percent per year, while the closed economies again grew at 0.7 percent per year. These findings are consistent with Smith's view that trade enhances prosperity, but they are not conclusive. Correlation does not prove causation. Perhaps being closed to trade is correlated with various other restrictive government policies, and it is those other policies that retard growth.

A second approach is to look at what happens when closed economies remove their trade restrictions. Once again, Smith's hypothesis fares well. Throughout history, when nations open themselves up to the world economy, the typical result is a subsequent increase in economic growth. This occurred in Japan in the 1850s, South Korea in the 1960s, and Vietnam in the 1990s. But once again, correlation does not prove causation. Trade liberalization is often accompanied by other reforms, and it is hard to disentangle the effects of trade from the effects of the other reforms.

A third approach to measuring the impact of trade on growth, proposed by economists Jeffrey Frankel and David Romer, is to look at the impact of geography. Some countries trade less simply because they are geographically disadvantaged. For example, New Zealand is disadvantaged compared to Belgium because it is farther from other populous countries. Similarly, landlocked countries are disadvantaged compared to countries with their own seaports. Because these geographical characteristics are correlated with trade, but arguably uncorrelated with other determinants of economic prosperity, they can be used to identify the causal impact of trade on income. (The statistical technique, which you may have studied in an econometrics course, is called *instrumental variables*.) After analyzing the data, Frankel and Romer conclude that "a rise of one percentage point in the ratio of trade to GDP increases income per person by at least one-half percentage point. Trade appears to raise income by spurring the accumulation of human and physical capital and by increasing output for given levels of capital."

The overwhelming weight of the evidence from this body of research is that Adam Smith was right. Openness to international trade is good for economic growth.⁴ ■

8-3 Policies to Promote Growth

So far we have used the Solow model to uncover the theoretical relationships among the different sources of economic growth, and we have discussed some of the empirical work that describes actual growth experiences. We can now use the theory and evidence to help guide our thinking about economic policy.

⁴ Jeffrey D. Sachs and Andrew Warner, "Economic Reform and the Process of Global Integration," *Brookings Papers on Economic Activity* (1995): 1–95; and Jeffrey A. Frankel and David Romer, "Does Trade Cause Growth?" *American Economic Review* 89 (June 1999): 379–399.

Evaluating the Rate of Saving

According to the Solow growth model, how much a nation saves and invests is a key determinant of its citizens' standard of living. So let's begin our policy discussion with a natural question: is the rate of saving in the U.S. economy too low, too high, or about right?

As we have seen, the saving rate determines the steady-state levels of capital and output. One particular saving rate produces the Golden Rule steady state, which maximizes consumption per worker and thus economic well-being. The Golden Rule provides the benchmark against which we can compare the U.S. economy.

To decide whether the U.S. economy is at, above, or below the Golden Rule steady state, we need to compare the marginal product of capital net of depreciation ($MPK - \delta$) with the growth rate of total output ($n + g$). As we established in Section 8-1, at the Golden Rule steady state, $MPK - \delta = n + g$. If the economy is operating with less capital than in the Golden Rule steady state, then diminishing marginal product tells us that $MPK - \delta > n + g$. In this case, increasing the rate of saving will increase capital accumulation and economic growth and, eventually, lead to a steady state with higher consumption (although consumption will be lower for part of the transition to the new steady state). On the other hand, if the economy has more capital than in the Golden Rule steady state, then $MPK - \delta < n + g$. In this case, capital accumulation is excessive: reducing the rate of saving will lead to higher consumption both immediately and in the long run.

To make this comparison for a real economy, such as the U.S. economy, we need an estimate of the growth rate of output ($n + g$) and an estimate of the net marginal product of capital ($MPK - \delta$). Real GDP in the United States grows an average of 3 percent per year, so $n + g = 0.03$. We can estimate the net marginal product of capital from the following three facts:

1. The capital stock is about 2.5 times one year's GDP.
2. Depreciation of capital is about 10 percent of GDP.
3. Capital income is about 30 percent of GDP.

Using the notation of our model (and the result from Chapter 3 that capital owners earn income of MPK for each unit of capital), we can write these facts as

1. $k = 2.5y$.
2. $\delta k = 0.1y$.
3. $MPK \times k = 0.3y$.

We solve for the rate of depreciation δ by dividing equation 2 by equation 1:

$$\begin{aligned}\delta k/k &= (0.1y)/(2.5y) \\ \delta &= 0.04.\end{aligned}$$

And we solve for the marginal product of capital MPK by dividing equation 3 by equation 1:

$$\begin{aligned}(MPK \times k)/k &= (0.3y)/(2.5y) \\ MPK &= 0.12.\end{aligned}$$

Thus, about 4 percent of the capital stock depreciates each year, and the marginal product of capital is about 12 percent per year. The net marginal product of capital, $MPK - \delta$, is about 8 percent per year.

We can now see that the return to capital ($MPK - \delta = 8$ percent per year) is well in excess of the economy's average growth rate ($n + g = 3$ percent per year). This fact, together with our previous analysis, indicates that the capital stock in the U.S. economy is well below the Golden Rule level. In other words, if the United States saved and invested a higher fraction of its income, it would grow more rapidly and eventually reach a steady state with higher consumption.

This conclusion is not unique to the U.S. economy. When calculations similar to those above are done for other economies, the results are similar. The possibility of excessive saving and capital accumulation beyond the Golden Rule level is intriguing as a matter of theory, but it appears not to be a problem that actual economies face. In practice, economists are more often concerned with insufficient saving. It is this kind of calculation that provides the intellectual foundation for this concern.⁵

Changing the Rate of Saving

The preceding calculations show that to move the U.S. economy toward the Golden Rule steady state, policymakers should increase national saving. But how can they do that? We saw in Chapter 3 that, as a matter of sheer accounting, higher national saving means higher public saving, higher private saving, or some combination of the two. Much of the debate over policies to increase growth centers on which of these options is likely to be most effective.

The most direct way in which the government affects national saving is through public saving—the difference between what the government receives in tax revenue and what it spends. When its spending exceeds its revenue, the government runs a *budget deficit*, which represents negative public saving. As we saw in Chapter 3, a budget deficit raises interest rates and crowds out investment; the resulting reduction in the capital stock is part of the burden of the national debt on future generations. Conversely, if it spends less than it raises in revenue, the government runs a *budget surplus*, which it can use to retire some of the national debt and stimulate investment.

The government also affects national saving by influencing private saving—the saving done by households and firms. In particular, how much people decide to save depends on the incentives they face, and these incentives are altered by a variety of public policies. Many economists argue that high tax rates on capital—including the corporate income tax, the federal income tax, the estate tax, and many state income and estate taxes—discourage private saving by reducing the rate of return that savers earn. On the other hand, tax-exempt retirement accounts, such as IRAs, are designed to encourage private saving by

⁵ For more on this topic and some international evidence, see Andrew B. Abel, N. Gregory Mankiw, Lawrence H. Summers, and Richard J. Zeckhauser, “Assessing Dynamic Efficiency: Theory and Evidence,” *Review of Economic Studies* 56 (1989): 1–19.

giving preferential treatment to income saved in these accounts. Some economists have proposed increasing the incentive to save by replacing the current system of income taxation with a system of consumption taxation.

Many disagreements over public policy are rooted in different views about how much private saving responds to incentives. For example, suppose that the government were to increase the amount that people can put into tax-exempt retirement accounts. Would people respond to this incentive by saving more? Or, instead, would people merely transfer saving already done in other forms into these accounts—reducing tax revenue and thus public saving without any stimulus to private saving? The desirability of the policy depends on the answers to these questions. Unfortunately, despite much research on this issue, no consensus has emerged.

Allocating the Economy's Investment

The Solow model makes the simplifying assumption that there is only one type of capital. In the world, of course, there are many types. Private businesses invest in traditional types of capital, such as bulldozers and steel plants, and newer types of capital, such as computers and robots. The government invests in various forms of public capital, called *infrastructure*, such as roads, bridges, and sewer systems.

In addition, there is *human capital*—the knowledge and skills that workers acquire through education, from early childhood programs such as Head Start to on-the-job training for adults in the labor force. Although the capital variable in the Solow model is usually interpreted as including only physical capital, in many ways human capital is analogous to physical capital. Like physical capital, human capital increases our ability to produce goods and services. Raising the level of human capital requires investment in the form of teachers, libraries, and student time. Recent research on economic growth has emphasized that human capital is at least as important as physical capital in explaining international differences in standards of living. One way of modeling this fact is to give the variable we call “capital” a broader definition that includes both human and physical capital.⁶

Policymakers trying to stimulate economic growth must confront the issue of what kinds of capital the economy needs most. In other words, what kinds of capital yield the highest marginal products? To a large extent, policymakers can rely on the marketplace to allocate the pool of saving to alternative types of investment. Those industries with the highest marginal products of capital will

⁶ Earlier in this chapter, when we were interpreting K as only physical capital, human capital was folded into the efficiency-of-labor parameter E . The alternative approach suggested here is to include human capital as part of K instead, so E represents technology but not human capital. If K is given this broader interpretation, then much of what we call labor income is really the return to human capital. As a result, the true capital share is much larger than the traditional Cobb–Douglas value of about 1/3. For more on this topic, see N. Gregory Mankiw, David Romer, and David N. Weil, “A Contribution to the Empirics of Economic Growth,” *Quarterly Journal of Economics* (May 1992): 407–437.

naturally be most willing to borrow at market interest rates to finance new investment. Many economists advocate that the government should merely create a “level playing field” for different types of capital—for example, by ensuring that the tax system treats all forms of capital equally. The government can then rely on the market to allocate capital efficiently.

Other economists have suggested that the government should actively encourage particular forms of capital. Suppose, for instance, that technological advance occurs as a by-product of certain economic activities. This would happen if new and improved production processes are devised during the process of building capital (a phenomenon called *learning by doing*) and if these ideas become part of society’s pool of knowledge. Such a by-product is called a *technological externality* (or a *knowledge spillover*). In the presence of such externalities, the social returns to capital exceed the private returns, and the benefits of increased capital accumulation to society are greater than the Solow model suggests.⁷ Moreover, some types of capital accumulation may yield greater externalities than others. If, for example, installing robots yields greater technological externalities than building a new steel mill, then perhaps the government should use the tax laws to encourage investment in robots. The success of such an *industrial policy*, as it is sometimes called, requires that the government be able to measure accurately the externalities of different economic activities so it can give the correct incentive to each activity.

Most economists are skeptical about industrial policies for two reasons. First, measuring the externalities from different sectors is virtually impossible. If policy is based on poor measurements, its effects might be close to random and, thus, worse than no policy at all. Second, the political process is far from perfect. Once the government gets into the business of rewarding specific industries with subsidies and tax breaks, the rewards are as likely to be based on political clout as on the magnitude of externalities.

One type of capital that necessarily involves the government is public capital. Local, state, and federal governments are always deciding if and when they should borrow to finance new roads, bridges, and transit systems. In 2009, one of President Barack Obama’s first economic proposals was to increase spending on such infrastructure. This policy was motivated by a desire partly to increase short-run aggregate demand (a goal we will examine later in this book) and partly to provide public capital and enhance long-run economic growth. Among economists, this policy had both defenders and critics. Yet all of them agree that measuring the marginal product of public capital is difficult. Private capital generates an easily measured rate of profit for the firm owning the capital, whereas the benefits of public capital are more diffuse. Furthermore, while private capital investment is made by investors spending their own money, the allocation of resources for public capital involves the political process and taxpayer funding. It is all too common to see “bridges to nowhere” being built simply because the local senator or congressman has the political muscle to get funds approved.

⁷ Paul Romer, “Crazy Explanations for the Productivity Slowdown,” *NBER Macroeconomics Annual* 2 (1987): 163–201.

Establishing the Right Institutions

As we discussed earlier, economists who study international differences in the standard of living attribute some of these differences to the inputs of physical and human capital and some to the productivity with which these inputs are used. One reason nations may have different levels of production efficiency is that they have different institutions guiding the allocation of scarce resources. Creating the right institutions is important for ensuring that resources are allocated to their best use.

A nation's legal tradition is an example of such an institution. Some countries, such as the United States, Australia, India, and Singapore, are former colonies of the United Kingdom and, therefore, have English-style common-law systems. Other nations, such as Italy, Spain, and most of those in Latin America, have legal traditions that evolved from the French Napoleonic Code. Studies have found that legal protections for shareholders and creditors are stronger in English-style than French-style legal systems. As a result, the English-style countries have better-developed capital markets. Nations with better-developed capital markets, in turn, experience more rapid growth because it is easier for small and start-up companies to finance investment projects, leading to a more efficient allocation of the nation's capital.⁸

Another important institutional difference across countries is the quality of government itself. Ideally, governments should provide a “helping hand” to the market system by protecting property rights, enforcing contracts, promoting competition, prosecuting fraud, and so on. Yet governments sometimes diverge from this ideal and act more like a “grabbing hand” by using the authority of the state to enrich a few powerful individuals at the expense of the broader community. Empirical studies have shown that the extent of corruption in a nation is indeed a significant determinant of economic growth.⁹

Adam Smith, the great eighteenth-century economist, was well aware of the role of institutions in economic growth. He once wrote, “Little else is requisite to carry a state to the highest degree of opulence from the lowest barbarism but peace, easy taxes, and a tolerable administration of justice: all the rest being brought about by the natural course of things.” Sadly, many nations do not enjoy these three simple advantages.

CASE STUDY

The Colonial Origins of Modern Institutions

International data show a remarkable correlation between latitude and economic prosperity: nations closer to the equator typically have lower levels of income per person than nations farther from the equator. This fact is true in both the northern and southern hemispheres.

⁸ Rafael La Porta, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert Vishny, “Law and Finance,” *Journal of Political Economy* 106 (1998): 1113–1155; and Ross Levine and Robert G. King, “Finance and Growth: Schumpeter Might Be Right,” *Quarterly Journal of Economics* 108 (1993): 717–737.

⁹ Paulo Mauro, “Corruption and Growth,” *Quarterly Journal of Economics* 110 (1995): 681–712.

What explains the correlation? Some economists have suggested that the tropical climates near the equator have a direct negative impact on productivity. In the heat of the tropics, agriculture is more difficult, and disease is more prevalent. This makes the production of goods and services more difficult.

Although the direct impact of geography is one reason tropical nations tend to be poor, it is not the whole story. Recent research by Daron Acemoglu, Simon Johnson, and James Robinson has suggested an indirect mechanism—the impact of geography on institutions. Here is their explanation, presented in several steps:

1. In the seventeenth, eighteenth, and nineteenth centuries, tropical climates presented European settlers with an increased risk of disease, especially malaria and yellow fever. As a result, when Europeans were colonizing much of the rest of the world, they avoided settling in tropical areas, such as most of Africa and Central America. The European settlers preferred areas with more moderate climates and better health conditions, such as the regions that are now the United States, Canada, and New Zealand.
2. In those areas where Europeans settled in large numbers, the settlers established European-like institutions that protected individual property rights and limited the power of government. By contrast, in tropical climates, the colonial powers often set up “extractive” institutions, including authoritarian governments, so they could take advantage of the area’s natural resources. These institutions enriched the colonizers, but they did little to foster economic growth.
3. Although the era of colonial rule is now long over, the early institutions that the European colonizers established are strongly correlated with the modern institutions in the former colonies. In tropical nations, where the colonial powers set up extractive institutions, there is typically less protection of property rights even today. When the colonizers left, the extractive institutions remained and were simply taken over by new ruling elites.
4. The quality of institutions is a key determinant of economic performance. Where property rights are well protected, people have more incentive to make the investments that lead to economic growth. Where property rights are less respected, as is typically the case in tropical nations, investment and growth tend to lag behind.

This research suggests that much of the international variation in living standards that we observe today is a result of the long reach of history.¹⁰ ■

Encouraging Technological Progress

The Solow model shows that sustained growth in income per worker must come from technological progress. The Solow model, however, takes technological progress as exogenous; it does not explain it. Unfortunately, the determinants of technological progress are not well understood.

¹⁰ Daron Acemoglu, Simon Johnson, and James A. Robinson, “The Colonial Origins of Comparative Development: An Empirical Investigation,” *American Economic Association* 91 (December 2001): 1369–1401.

Despite this limited understanding, many public policies are designed to stimulate technological progress. Most of these policies encourage the private sector to devote resources to technological innovation. For example, the patent system gives a temporary monopoly to inventors of new products; the tax code offers tax breaks for firms engaging in research and development; and government agencies, such as the National Science Foundation, directly subsidize basic research in universities. In addition, as discussed above, proponents of industrial policy argue that the government should take a more active role in promoting specific industries that are key for rapid technological advance.

In recent years, the encouragement of technological progress has taken on an international dimension. Many of the companies that engage in research to advance technology are located in the United States and other developed nations. Developing nations such as China have an incentive to “free ride” on this research by not strictly enforcing intellectual property rights. That is, Chinese companies often use the ideas developed abroad without compensating the patent holders. The United States has strenuously objected to this practice, and China has promised to step up enforcement. If intellectual property rights were better enforced around the world, firms would have more incentive to engage in research, and this would promote worldwide technological progress.

CASE STUDY

The Worldwide Slowdown in Economic Growth: 1972–1995

Beginning in the early 1970s, and lasting until the mid-1990s, world policymakers faced a perplexing problem: a global slowdown in economic growth. Table 8-2 presents data on the growth in real GDP per person for the seven major economies. Growth in the United States fell from 2.2 percent before 1972 to 1.5 percent from 1972 to 1995. Other countries experienced similar or more severe declines. Accumulated over many years, even a small change in the rate of growth has a large effect on economic well-being. Real income in the United States today is almost 20 percent lower than it would have been had growth remained at its previous level.

Why did this slowdown occur? Studies have shown that it was attributable to a fall in the rate at which the production function was improving over time. The appendix to this chapter explains how economists measure changes in the production function with a variable called *total factor productivity*, which is closely related to the efficiency of labor in the Solow model. There are many hypotheses to explain this fall in productivity growth. Here are four of them.

Measurement Problems One possibility is that the productivity slowdown did not really occur and that it shows up in the data because the data are flawed. As you may recall from Chapter 2, one problem in measuring inflation is correcting for changes in the quality of goods and services. The same issue arises when measuring output and productivity. For instance, if technological advance leads to *more* computers being built, then the increase in output and productivity

TABLE 8-2

Growth Around the World

Country	GROWTH IN OUTPUT PER PERSON (PERCENT PER YEAR)		
	1948–1972	1972–1995	1995–2007
Canada	2.9	1.8	2.2
France	4.3	1.6	1.7
West Germany	5.7	2.0	
Germany			1.5
Italy	4.9	2.3	1.2
Japan	8.2	2.6	1.2
United Kingdom	2.4	1.8	2.6
United States	2.2	1.5	2.0

Source: Angus Maddison, *Phases of Capitalist Development* (Oxford: Oxford University Press, 1982); OECD National Accounts; and World Bank: *World Development Indicators*.

is easy to measure. But if technological advance leads to *faster* computers being built, then output and productivity have increased, but that increase is more subtle and harder to measure. Government statisticians try to correct for changes in quality, but despite their best efforts, the resulting data are far from perfect.

Unmeasured quality improvements mean that our standard of living is rising more rapidly than the official data indicate. This issue should make us suspicious of the data, but by itself it cannot explain the productivity slowdown. To explain a *slowdown* in growth, one must argue that the measurement problems got *worse*. There is some indication that this might be so. As history passes, fewer people work in industries with tangible and easily measured output, such as agriculture, and more work in industries with intangible and less easily measured output, such as medical services. Yet few economists believe that measurement problems were the full story.

Oil Prices When the productivity slowdown began around 1973, the obvious hypothesis to explain it was the large increase in oil prices caused by the actions of the OPEC oil cartel. The primary piece of evidence was the timing: productivity growth slowed at the same time that oil prices skyrocketed. Over time, however, this explanation has appeared less likely. One reason is that the accumulated shortfall in productivity seems too large to be explained by an increase in oil prices—petroleum-based products are not that large a fraction of a typical firm's costs. In addition, if this explanation were right, productivity should have sped up when political turmoil in OPEC caused oil prices to plummet in 1986. Unfortunately, that did not happen.

Worker Quality Some economists suggest that the productivity slowdown might have been caused by changes in the labor force. In the early 1970s, the large baby-boom generation started leaving school and taking jobs. At the same

time, changing social norms encouraged many women to leave full-time housework and enter the labor force. Both of these developments lowered the average level of experience among workers, which in turn lowered average productivity.

Other economists point to changes in worker quality as gauged by human capital. Although the educational attainment of the labor force continued to rise throughout this period, it was not increasing as rapidly as it had in the past. Moreover, declining performance on some standardized tests suggests that the quality of education was declining. If so, this could explain slowing productivity growth.

The Depletion of Ideas Still other economists suggest that the world started to run out of new ideas about how to produce in the early 1970s, pushing the economy into an age of slower technological progress. These economists often argue that the anomaly is not the period since 1970 but the preceding two decades. In the late 1940s, the economy had a large backlog of ideas that had not been fully implemented because of the Great Depression of the 1930s and World War II in the first half of the 1940s. After the economy used up this backlog, the argument goes, a slowdown in productivity growth was likely. Indeed, although the growth rates in the 1970s, 1980s, and early 1990s were disappointing compared to those of the 1950s and 1960s, they were not lower than average growth rates from 1870 to 1950.

As any good doctor will tell you, sometimes a patient's illness goes away on its own, even if the doctor has failed to come up with a convincing diagnosis and remedy. This seems to be the outcome of the productivity slowdown. In the middle of the 1990s, economic growth took off, at least in the English-speaking countries of the United States, Canada, and the United Kingdom. As with the slowdown in economic growth in the 1970s, the acceleration in the 1990s is hard to explain definitively. But at least part of the credit goes to advances in computer and information technology, including the Internet.¹¹ ■

8-4 Beyond the Solow Model: Endogenous Growth Theory

A chemist, a physicist, and an economist are all trapped on a desert island, trying to figure out how to open a can of food.

“Let’s heat the can over the fire until it explodes,” says the chemist.

“No, no,” says the physicist, “let’s drop the can onto the rocks from the top of a high tree.”

“I have an idea,” says the economist. “First, we assume a can opener . . .”

¹¹ For various views on the growth slowdown, see “Symposium: The Slowdown in Productivity Growth” in the Fall 1988 issue of *The Journal of Economic Perspectives*. For a discussion of the subsequent growth acceleration and the role of information technology, see “Symposium: Computers and Productivity” in the Fall 2000 issue of *The Journal of Economic Perspectives*.

This old joke takes aim at how economists use assumptions to simplify—and sometimes oversimplify—the problems they face. It is particularly apt when evaluating the theory of economic growth. One goal of growth theory is to explain the persistent rise in living standards that we observe in most parts of the world. The Solow growth model shows that such persistent growth must come from technological progress. But where does technological progress come from? In the Solow model, it is just assumed!

The preceding Case Study on the productivity slowdown of the 1970s and speed-up of the 1990s suggests that changes in the pace of technological progress are tremendously important. To understand fully the process of economic growth, we need to go beyond the Solow model and develop models that explain technological advance. Models that do this often go by the label **endogenous growth theory** because they reject the Solow model's assumption of exogenous technological change. Although the field of endogenous growth theory is large and sometimes complex, here we get a quick taste of this modern research.¹²

The Basic Model

To illustrate the idea behind endogenous growth theory, let's start with a particularly simple production function:

$$Y = AK,$$

where Y is output, K is the capital stock, and A is a constant measuring the amount of output produced for each unit of capital. Notice that this production function does not exhibit the property of diminishing returns to capital. One extra unit of capital produces A extra units of output, regardless of how much capital there is. This absence of diminishing returns to capital is the key difference between this endogenous growth model and the Solow model.

Now let's see what this production function says about economic growth. As before, we assume a fraction s of income is saved and invested. We therefore describe capital accumulation with an equation similar to those we used previously:

$$\Delta K = sY - \delta K.$$

This equation states that the change in the capital stock (ΔK) equals investment (sY) minus depreciation (δK). Combining this equation with the $Y = AK$ production function, we obtain, after a bit of manipulation,

$$\Delta Y/Y = \Delta K/K = sA - \delta.$$

¹² This section provides a brief introduction to the large and fascinating literature on endogenous growth theory. Early and important contributions to this literature include Paul M. Romer, "Increasing Returns and Long-Run Growth," *Journal of Political Economy* 94 (October 1986): 1002–1037; and Robert E. Lucas, Jr., "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22 (1988): 3–42. The reader can learn more about this topic in the undergraduate textbook by David N. Weil, *Economic Growth*, 2nd ed. (Pearson, 2008).

This equation shows what determines the growth rate of output $\Delta Y/Y$. Notice that, as long as $sA > \delta$, the economy's income grows forever, even without the assumption of exogenous technological progress.

Thus, a simple change in the production function can alter dramatically the predictions about economic growth. In the Solow model, saving leads to growth temporarily, but diminishing returns to capital eventually force the economy to approach a steady state in which growth depends only on exogenous technological progress. By contrast, in this endogenous growth model, saving and investment can lead to persistent growth.

But is it reasonable to abandon the assumption of diminishing returns to capital? The answer depends on how we interpret the variable K in the production function $Y = AK$. If we take the traditional view that K includes only the economy's stock of plants and equipment, then it is natural to assume diminishing returns. Giving 10 computers to a worker does not make that worker 10 times as productive as he or she is with one computer.

Advocates of endogenous growth theory, however, argue that the assumption of constant (rather than diminishing) returns to capital is more palatable if K is interpreted more broadly. Perhaps the best case can be made for the endogenous growth model by viewing knowledge as a type of capital. Clearly, knowledge is an important input into the economy's production—both its production of goods and services and its production of new knowledge. Compared to other forms of capital, however, it is less natural to assume that knowledge exhibits the property of diminishing returns. (Indeed, the increasing pace of scientific and technological innovation over the past few centuries has led some economists to argue that there are increasing returns to knowledge.) If we accept the view that knowledge is a type of capital, then this endogenous growth model with its assumption of constant returns to capital becomes a more plausible description of long-run economic growth.

A Two-Sector Model

Although the $Y = AK$ model is the simplest example of endogenous growth, the theory has gone well beyond this. One line of research has tried to develop models with more than one sector of production in order to offer a better description of the forces that govern technological progress. To see what we might learn from such models, let's sketch out an example.

The economy has two sectors, which we can call manufacturing firms and research universities. Firms produce goods and services, which are used for consumption and investment in physical capital. Universities produce a factor of production called "knowledge," which is then freely used in both sectors. The economy is described by the production function for firms, the production function for universities, and the capital-accumulation equation:

$$\begin{aligned}
 Y &= F[K, (1 - u)LE] && \text{(production function in manufacturing firms),} \\
 \Delta E &= g(u)E && \text{(production function in research universities),} \\
 \Delta K &= sY - \delta K && \text{(capital accumulation),}
 \end{aligned}$$

where u is the fraction of the labor force in universities (and $1 - u$ is the fraction in manufacturing), E is the stock of knowledge (which in turn determines the efficiency of labor), and g is a function that shows how the growth in knowledge depends on the fraction of the labor force in universities. The rest of the notation is standard. As usual, the production function for the manufacturing firms is assumed to have constant returns to scale: if we double both the amount of physical capital (K) and the effective number of workers in manufacturing $[(1 - u)LE]$, we double the output of goods and services (Y).

This model is a cousin of the $Y = AK$ model. Most important, this economy exhibits constant (rather than diminishing) returns to capital, as long as capital is broadly defined to include knowledge. In particular, if we double both physical capital K and knowledge E , then we double the output of both sectors in the economy. As a result, like the $Y = AK$ model, this model can generate persistent growth without the assumption of exogenous shifts in the production function. Here persistent growth arises endogenously because the creation of knowledge in universities never slows down.

At the same time, however, this model is also a cousin of the Solow growth model. If u , the fraction of the labor force in universities, is held constant, then the efficiency of labor E grows at the constant rate $g(u)$. This result of constant growth in the efficiency of labor at rate g is precisely the assumption made in the Solow model with technological progress. Moreover, the rest of the model—the manufacturing production function and the capital-accumulation equation—also resembles the rest of the Solow model. As a result, for any given value of u , this endogenous growth model works just like the Solow model.

There are two key decision variables in this model. As in the Solow model, the fraction of output used for saving and investment, s , determines the steady-state stock of physical capital. In addition, the fraction of labor in universities, u , determines the growth in the stock of knowledge. Both s and u affect the level of income, although only u affects the steady-state growth rate of income. Thus, this model of endogenous growth takes a small step in the direction of showing which societal decisions determine the rate of technological change.

The Microeconomics of Research and Development

The two-sector endogenous growth model just presented takes us closer to understanding technological progress, but it still tells only a rudimentary story about the creation of knowledge. If one thinks about the process of research and development for even a moment, three facts become apparent. First, although knowledge is largely a public good (that is, a good freely available to everyone), much research is done in firms that are driven by the profit motive. Second, research is profitable because innovations give firms temporary monopolies, either because of the patent system or because there is an advantage to being the first firm on the market with a new product. Third, when one firm innovates, other firms build on that innovation to produce the next generation of innovations. These (essentially microeconomic) facts are not easily connected with the (essentially macroeconomic) growth models we have discussed so far.

Some endogenous growth models try to incorporate these facts about research and development. Doing this requires modeling both the decisions that firms face as they engage in research and the interactions among firms that have some degree of monopoly power over their innovations. Going into more detail about these models is beyond the scope of this book, but it should be clear already that one virtue of these endogenous growth models is that they offer a more complete description of the process of technological innovation.

One question these models are designed to address is whether, from the standpoint of society as a whole, private profit-maximizing firms tend to engage in too little or too much research. In other words, is the social return to research (which is what society cares about) greater or smaller than the private return (which is what motivates individual firms)? It turns out that, as a theoretical matter, there are effects in both directions. On the one hand, when a firm creates a new technology, it makes other firms better off by giving them a base of knowledge on which to build in future research. As Isaac Newton famously remarked, “If I have seen farther than others, it is because I was standing on the shoulders of giants.” On the other hand, when one firm invests in research, it can also make other firms worse off if it does little more than being the first to discover a technology that another firm would have invented in due course. This duplication of research effort has been called the “stepping on toes” effect. Whether firms left to their own devices do too little or too much research depends on whether the positive “standing on shoulders” externality or the negative “stepping on toes” externality is more prevalent.

Although theory alone is ambiguous about whether research effort is more or less than optimal, the empirical work in this area is usually less so. Many studies have suggested the “standing on shoulders” externality is important and, as a result, the social return to research is large—often in excess of 40 percent per year. This is an impressive rate of return, especially when compared to the return to physical capital, which we earlier estimated to be about 8 percent per year. In the judgment of some economists, this finding justifies substantial government subsidies to research.¹³

The Process of Creative Destruction

In his 1942 book *Capitalism, Socialism, and Democracy*, economist Joseph Schumpeter suggested that economic progress comes through a process of “creative destruction.” According to Schumpeter, the driving force behind progress is the entrepreneur with an idea for a new product, a new way to produce an old product, or some other innovation. When the entrepreneur’s firm enters the market, it has some degree of monopoly power over its innovation; indeed, it is the prospect of monopoly profits that motivates the entrepreneur. The entry of the new firm is good for consumers, who now have an expanded range of choices, but it is often

¹³ For an overview of the empirical literature on the effects of research, see Zvi Griliches, “The Search for R&D Spillovers,” *Scandinavian Journal of Economics* 94 (1991): 29–47.

bad for incumbent producers, who may find it hard to compete with the entrant. If the new product is sufficiently better than old ones, the incumbents may even be driven out of business. Over time, the process keeps renewing itself. The entrepreneur's firm becomes an incumbent, enjoying high profitability until its product is displaced by another entrepreneur with the next generation of innovation.

History confirms Schumpeter's thesis that there are winners and losers from technological progress. For example, in England in the early nineteenth century, an important innovation was the invention and spread of machines that could produce textiles using unskilled workers at low cost. This technological advance was good for consumers, who could clothe themselves more cheaply. Yet skilled knitters in England saw their jobs threatened by new technology, and they responded by organizing violent revolts. The rioting workers, called Luddites, smashed the weaving machines used in the wool and cotton mills and set the homes of the mill owners on fire (a less than creative form of destruction). Today, the term "Luddite" refers to anyone who opposes technological progress.

A more recent example of creative destruction involves the retailing giant Wal-Mart. Although retailing may seem like a relatively static activity, in fact it is a sector that has seen sizable rates of technological progress over the past several decades. Through better inventory-control, marketing, and personnel-management techniques, for example, Wal-Mart has found ways to bring goods to consumers at lower cost than traditional retailers. These changes benefit consumers, who can buy goods at lower prices, and the stockholders of Wal-Mart, who share in its profitability. But they adversely affect small mom-and-pop stores, which find it hard to compete when a Wal-Mart opens nearby.

Faced with the prospect of being the victims of creative destruction, incumbent producers often look to the political process to stop the entry of new, more efficient competitors. The original Luddites wanted the British government to save their jobs by restricting the spread of the new textile technology; instead, Parliament sent troops to suppress the Luddite riots. Similarly, in recent years, local retailers have sometimes tried to use local land-use regulations to stop Wal-Mart from entering their market. The cost of such entry restrictions, however, is to slow the pace of technological progress. In Europe, where entry regulations are stricter than they are in the United States, the economies have not seen the emergence of retailing giants like Wal-Mart; as a result, productivity growth in retailing has been much lower.¹⁴

Schumpeter's vision of how capitalist economies work has merit as a matter of economic history. Moreover, it has inspired some recent work in the theory of economic growth. One line of endogenous growth theory, pioneered by economists Philippe Aghion and Peter Howitt, builds on Schumpeter's insights by modeling technological advance as a process of entrepreneurial innovation and creative destruction.¹⁵

¹⁴ Robert J. Gordon, "Why Was Europe Left at the Station When America's Productivity Locomotive Departed?" *NBER Working Paper* No. 10661, 2004.

¹⁵ Philippe Aghion and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica* 60 (1992): 323–351.

8-5 Conclusion

Long-run economic growth is the single most important determinant of the economic well-being of a nation's citizens. Everything else that macroeconomists study—unemployment, inflation, trade deficits, and so on—pales in comparison.

Fortunately, economists know quite a lot about the forces that govern economic growth. The Solow growth model and the more recent endogenous growth models show how saving, population growth, and technological progress interact in determining the level and growth of a nation's standard of living. These theories offer no magic recipe to ensure an economy achieves rapid growth, but they give much insight, and they provide the intellectual framework for much of the debate over public policy aimed at promoting long-run economic growth.

Summary

1. In the steady state of the Solow growth model, the growth rate of income per person is determined solely by the exogenous rate of technological progress.
2. Many empirical studies have examined to what extent the Solow model can help explain long-run economic growth. The model can explain much of what we see in the data, such as balanced growth and conditional convergence. Recent studies have also found that international variation in standards of living is attributable to a combination of capital accumulation and the efficiency with which capital is used.
3. In the Solow model with population growth and technological progress, the Golden Rule (consumption-maximizing) steady state is characterized by equality between the net marginal product of capital ($MPK - \delta$) and the steady-state growth rate of total income ($n + g$). In the U.S. economy, the net marginal product of capital is well in excess of the growth rate, indicating that the U.S. economy has a lower saving rate and less capital than it would have in the Golden Rule steady state.
4. Policymakers in the United States and other countries often claim that their nations should devote a larger percentage of their output to saving and investment. Increased public saving and tax incentives for private saving are two ways to encourage capital accumulation. Policymakers can also promote economic growth by setting up the right legal and financial institutions so that resources are allocated efficiently and by ensuring proper incentives to encourage research and technological progress.
5. In the early 1970s, the rate of growth of income per person fell substantially in most industrialized countries, including the United States. The cause of this slowdown is not well understood. In the mid-1990s, the U.S. growth rate increased, most likely because of advances in information technology.

6. Modern theories of endogenous growth attempt to explain the rate of technological progress, which the Solow model takes as exogenous. These models try to explain the decisions that determine the creation of knowledge through research and development.

KEY CONCEPTS

Efficiency of labor

Labor-augmenting technological progress

Endogenous growth theory

QUESTIONS FOR REVIEW

- In the Solow model, what determines the steady-state rate of growth of income per worker?
- In the steady state of the Solow model, at what rate does output per person grow? At what rate does capital per person grow? How does this compare with the U.S. experience?
- What data would you need to determine whether an economy has more or less capital than in the Golden Rule steady state?
- How can policymakers influence a nation's saving rate?
- What has happened to the rate of productivity growth over the past 50 years? How might you explain this phenomenon?
- How does endogenous growth theory explain persistent growth without the assumption of exogenous technological progress? How does this differ from the Solow model?

PROBLEMS AND APPLICATIONS

1. An economy described by the Solow growth model has the following production function:

$$y = \sqrt{k}.$$

- Solve for the steady-state value of y as a function of s , n , g , and δ .
 - A developed country has a saving rate of 28 percent and a population growth rate of 1 percent per year. A less developed country has a saving rate of 10 percent and a population growth rate of 4 percent per year. In both countries, $g = 0.02$ and $\delta = 0.04$. Find the steady-state value of y for each country.
 - What policies might the less developed country pursue to raise its level of income?
2. In the United States, the capital share of GDP is about 30 percent, the average growth in output is about 3 percent per year, the depreciation rate is about 4 percent per year, and the capital-output ratio is about 2.5. Suppose that the production function is Cobb–Douglas, so

that the capital share in output is constant, and that the United States has been in a steady state. (For a discussion of the Cobb–Douglas production function, see Chapter 3.)

- What must the saving rate be in the initial steady state? [*Hint:* Use the steady-state relationship, $sy = (\delta + n + g)k$.]
- What is the marginal product of capital in the initial steady state?
- Suppose that public policy raises the saving rate so that the economy reaches the Golden Rule level of capital. What will the marginal product of capital be at the Golden Rule steady state? Compare the marginal product at the Golden Rule steady state to the marginal product in the initial steady state. Explain.
- What will the capital-output ratio be at the Golden Rule steady state? (*Hint:* For the Cobb–Douglas production function, the capital-output ratio is related to the marginal product of capital.)

- e. What must the saving rate be to reach the Golden Rule steady state?
3. Prove each of the following statements about the steady state of the Solow model with population growth and technological progress.
- The capital–output ratio is constant.
 - Capital and labor each earn a constant share of an economy’s income. [*Hint*: Recall the definition $MPK = f(k + 1) - f(k)$.]
 - Total capital income and total labor income both grow at the rate of population growth plus the rate of technological progress, $n + g$.
 - The real rental price of capital is constant, and the real wage grows at the rate of technological progress g . (*Hint*: The real rental price of capital equals total capital income divided by the capital stock, and the real wage equals total labor income divided by the labor force.)
4. Two countries, Richland and Poorland, are described by the Solow growth model. They have the same Cobb–Douglas production function, $F(K, L) = A K^\alpha L^{1-\alpha}$, but with different quantities of capital and labor. Richland saves 32 percent of its income, while Poorland saves 10 percent. Richland has population growth of 1 percent per year, while Poorland has population growth of 3 percent. (The numbers in this problem are chosen to be approximately realistic descriptions of rich and poor nations.) Both nations have technological progress at a rate of 2 percent per year and depreciation at a rate of 5 percent per year.
- What is the per-worker production function $f(k)$?
 - Solve for the ratio of Richland’s steady-state income per worker to Poorland’s. (*Hint*: The parameter α will play a role in your answer.)
 - If the Cobb–Douglas parameter α takes the conventional value of about $1/3$, how much higher should income per worker be in Richland compared to Poorland?
 - Income per worker in Richland is actually 16 times income per worker in Poorland. Can you explain this fact by changing the value of the parameter α ? What must it be? Can you think of any way of justifying such a value for this parameter? How else might you explain
- the large difference in income between Richland and Poorland?
5. The amount of education the typical person receives varies substantially among countries. Suppose you were to compare a country with a highly educated labor force and a country with a less educated labor force. Assume that education affects only the level of the efficiency of labor. Also assume that the countries are otherwise the same: they have the same saving rate, the same depreciation rate, the same population growth rate, and the same rate of technological progress. Both countries are described by the Solow model and are in their steady states. What would you predict for the following variables?
- The rate of growth of total income.
 - The level of income per worker.
 - The real rental price of capital.
 - The real wage.
6. This question asks you to analyze in more detail the two-sector endogenous growth model presented in the text.
- Rewrite the production function for manufactured goods in terms of output per effective worker and capital per effective worker.
 - In this economy, what is break-even investment (the amount of investment needed to keep capital per effective worker constant)?
 - Write down the equation of motion for k , which shows Δk as saving minus break-even investment. Use this equation to draw a graph showing the determination of steady-state k . (*Hint*: This graph will look much like those we used to analyze the Solow model.)
 - In this economy, what is the steady-state growth rate of output per worker Y/L ? How do the saving rate s and the fraction of the labor force in universities u affect this steady-state growth rate?
 - Using your graph, show the impact of an increase in u . (*Hint*: This change affects both curves.) Describe both the immediate and the steady-state effects.
 - Based on your analysis, is an increase in u an unambiguously good thing for the economy? Explain.



Accounting for the Sources of Economic Growth

Real GDP in the United States has grown an average of about 3 percent per year over the past 50 years. What explains this growth? In Chapter 3 we linked the output of the economy to the factors of production—capital and labor—and to the production technology. Here we develop a technique called *growth accounting* that divides the growth in output into three different sources: increases in capital, increases in labor, and advances in technology. This breakdown provides us with a measure of the rate of technological change.

Increases in the Factors of Production

We first examine how increases in the factors of production contribute to increases in output. To do this, we start by assuming there is no technological change, so the production function relating output Y to capital K and labor L is constant over time:

$$Y = F(K, L).$$

In this case, the amount of output changes only because the amount of capital or labor changes.

Increases in Capital First, consider changes in capital. If the amount of capital increases by ΔK units, by how much does the amount of output increase? To answer this question, we need to recall the definition of the marginal product of capital MPK :

$$MPK = F(K + 1, L) - F(K, L).$$

The marginal product of capital tells us how much output increases when capital increases by 1 unit. Therefore, when capital increases by ΔK units, output increases by approximately $MPK \times \Delta K$.¹⁶

For example, suppose that the marginal product of capital is $1/5$; that is, an additional unit of capital increases the amount of output produced by one-fifth

¹⁶ Note the word “approximately” here. This answer is only an approximation because the marginal product of capital varies: it falls as the amount of capital increases. An exact answer would take into account that each unit of capital has a different marginal product. If the change in K is not too large, however, the approximation of a constant marginal product is very accurate.

of a unit. If we increase the amount of capital by 10 units, we can compute the amount of additional output as follows:

$$\begin{aligned}\Delta Y &= \quad MPK \quad \times \quad \Delta K \\ &= 1/5 \frac{\text{units of output}}{\text{unit of capital}} \times 10 \text{ units of capital} \\ &= 2 \text{ units of output.}\end{aligned}$$

By increasing capital by 10 units, we obtain 2 more units of output. Thus, we use the marginal product of capital to convert changes in capital into changes in output.

Increases in Labor Next, consider changes in labor. If the amount of labor increases by ΔL units, by how much does output increase? We answer this question the same way we answered the question about capital. The marginal product of labor MPL tells us how much output changes when labor increases by 1 unit—that is,

$$MPL = F(K, L + 1) - F(K, L).$$

Therefore, when the amount of labor increases by ΔL units, output increases by approximately $MPL \times \Delta L$.

For example, suppose that the marginal product of labor is 2; that is, an additional unit of labor increases the amount of output produced by 2 units. If we increase the amount of labor by 10 units, we can compute the amount of additional output as follows:

$$\begin{aligned}\Delta Y &= \quad MPL \quad \times \quad \Delta L \\ &= 2 \frac{\text{units of output}}{\text{unit of labor}} \times 10 \text{ units of labor} \\ &= 20 \text{ units of output.}\end{aligned}$$

By increasing labor by 10 units, we obtain 20 more units of output. Thus, we use the marginal product of labor to convert changes in labor into changes in output.

Increases in Capital and Labor Finally, let's consider the more realistic case in which both factors of production change. Suppose that the amount of capital increases by ΔK and the amount of labor increases by ΔL . The increase in output then comes from two sources: more capital and more labor. We can divide this increase into the two sources using the marginal products of the two inputs:

$$\Delta Y = (MPK \times \Delta K) + (MPL \times \Delta L).$$

The first term in parentheses is the increase in output resulting from the increase in capital, and the second term in parentheses is the increase in output resulting from the increase in labor. This equation shows us how to attribute growth to each factor of production.

We now want to convert this last equation into a form that is easier to interpret and apply to the available data. First, with some algebraic rearrangement, the equation becomes¹⁷

$$\frac{\Delta Y}{Y} = \left(\frac{MPK \times K}{Y} \right) \frac{\Delta K}{K} + \left(\frac{MPL \times L}{Y} \right) \frac{\Delta L}{L}.$$

This form of the equation relates the growth rate of output, $\Delta Y/Y$, to the growth rate of capital, $\Delta K/K$, and the growth rate of labor, $\Delta L/L$.

Next, we need to find some way to measure the terms in parentheses in the last equation. In Chapter 3 we showed that the marginal product of capital equals its real rental price. Therefore, $MPK \times K$ is the total return to capital, and $(MPK \times K)/Y$ is capital's share of output. Similarly, the marginal product of labor equals the real wage. Therefore, $MPL \times L$ is the total compensation that labor receives, and $(MPL \times L)/Y$ is labor's share of output. Under the assumption that the production function has constant returns to scale, Euler's theorem (which we discussed in Chapter 3) tells us that these two shares sum to 1. In this case, we can write

$$\frac{\Delta Y}{Y} = \alpha \frac{\Delta K}{K} + (1 - \alpha) \frac{\Delta L}{L},$$

where α is capital's share and $(1 - \alpha)$ is labor's share.

This last equation gives us a simple formula for showing how changes in inputs lead to changes in output. It shows, in particular, that we must weight the growth rates of the inputs by the factor shares. As we discussed in Chapter 3, capital's share in the United States is about 30 percent, that is, $\alpha = 0.30$. Therefore, a 10-percent increase in the amount of capital ($\Delta K/K = 0.10$) leads to a 3-percent increase in the amount of output ($\Delta Y/Y = 0.03$). Similarly, a 10-percent increase in the amount of labor ($\Delta L/L = 0.10$) leads to a 7-percent increase in the amount of output ($\Delta Y/Y = 0.07$).

Technological Progress

So far in our analysis of the sources of growth, we have been assuming that the production function does not change over time. In practice, of course, technological progress improves the production function. For any given amount of inputs, we can produce more output today than we could in the past. We now extend the analysis to allow for technological progress.

We include the effects of the changing technology by writing the production function as

$$Y = AF(K, L),$$

¹⁷ *Mathematical note:* To see that this is equivalent to the previous equation, note that we can multiply both sides of this equation by Y and thereby cancel Y from three places in which it appears. We can cancel the K in the top and bottom of the first term on the right-hand side and the L in the top and bottom of the second term on the right-hand side. These algebraic manipulations turn this equation into the previous one.

where A is a measure of the current level of technology called *total factor productivity*. Output now increases not only because of increases in capital and labor but also because of increases in total factor productivity. If total factor productivity increases by 1 percent and if the inputs are unchanged, then output increases by 1 percent.

Allowing for a changing level of technology adds another term to our equation accounting for economic growth:

$$\frac{\Delta Y}{Y} = \alpha \frac{\Delta K}{K} + (1 - \alpha) \frac{\Delta L}{L} + \frac{\Delta A}{A}$$

Growth in Output = Contribution of Capital + Contribution of Labor + Growth in Total Factor Productivity

This is the key equation of growth accounting. It identifies and allows us to measure the three sources of growth: changes in the amount of capital, changes in the amount of labor, and changes in total factor productivity.

Because total factor productivity is not directly observable, it is measured indirectly. We have data on the growth in output, capital, and labor; we also have data on capital's share of output. From these data and the growth-accounting equation, we can compute the growth in total factor productivity to make sure that everything adds up:

$$\frac{\Delta A}{A} = \frac{\Delta Y}{Y} - \alpha \frac{\Delta K}{K} - (1 - \alpha) \frac{\Delta L}{L}$$

$\Delta A/A$ is the change in output that cannot be explained by changes in inputs. Thus, the growth in total factor productivity is computed as a residual—that is, as the amount of output growth that remains after we have accounted for the determinants of growth that we can measure directly. Indeed, $\Delta A/A$ is sometimes called the *Solow residual*, after Robert Solow, who first showed how to compute it.¹⁸

Total factor productivity can change for many reasons. Changes most often arise because of increased knowledge about production methods, so the Solow residual is often used as a measure of technological progress. Yet other factors, such as education and government regulation, can affect total factor productivity as well. For example, if higher public spending raises the quality of education, then workers may become more productive and output may rise, which implies higher total factor productivity. As another example, if government regulations require firms to purchase capital to reduce pollution or increase worker safety, then the capital stock may rise without any increase in measured output, which implies lower total factor productivity. *Total factor productivity captures anything that changes the relation between measured inputs and measured output.*

¹⁸ Robert M. Solow, "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics* 39 (1957): 312–320. It is natural to ask how growth in labor efficiency E relates to growth in total factor productivity. One can show that $\Delta A/A = (1 - \alpha)\Delta E/E$, where α is capital's share. Thus, technological change as measured by growth in the efficiency of labor is proportional to technological change as measured by the Solow residual.

The Sources of Growth in the United States

Having learned how to measure the sources of economic growth, we now look at the data. Table 8-3 uses U.S. data to measure the contributions of the three sources of growth between 1948 and 2007.

This table shows that output in the non-farm business sector grew an average of 3.6 percent per year during this time. Of this 3.6 percent, 1.2 percent was attributable to increases in the capital stock, 1.2 percent to increases in the labor input, and 1.2 percent to increases in total factor productivity. These data show that increases in capital, labor, and productivity have contributed almost equally to economic growth in the United States.

Table 8-3 also shows that the growth in total factor productivity slowed substantially during the period from 1972 to 1995. In a case study in this chapter, we discussed some hypotheses to explain this productivity slowdown.

TABLE 8-3

Accounting for Economic Growth in the United States

Years	Output Growth $\Delta Y/Y$	SOURCE OF GROWTH		
		Capital $\alpha\Delta K/K$	Labor $(1 - \alpha)\Delta L/L$	Total Factor Productivity $\Delta A/A$
		(average percentage increase per year)		
1948–2007	3.6	1.2	1.2	1.2
1948–1972	4.0	1.2	0.9	1.9
1972–1995	3.4	1.3	1.5	0.6
1995–2007	3.5	1.3	1.0	1.3

Source: U.S. Department of Labor. Data are for the non-farm business sector.

CASE STUDY

Growth in the East Asian Tigers

Perhaps the most spectacular growth experiences in recent history have been those of the “Tigers” of East Asia: Hong Kong, Singapore, South Korea, and Taiwan. From 1966 to 1990, while real income per person was growing about 2 percent per year in the United States, it grew more than 7 percent per year in each of these countries. In the course of a single generation, real income per person increased fivefold, moving the Tigers from among the world’s poorest countries to among the richest. (In the late 1990s, a period of pronounced financial turmoil tarnished the reputation of some of these economies. But this short-run problem, which we examine in a case study in Chapter 12, doesn’t

come close to reversing the spectacular long-run growth that the Asian Tigers have experienced.)

What accounts for these growth miracles? Some commentators have argued that the success of these four countries is hard to reconcile with basic growth theory, such as the Solow growth model, which takes technology as growing at a constant, exogenous rate. They have suggested that these countries' rapid growth is explained by their ability to imitate foreign technologies. By adopting technology developed abroad, the argument goes, these countries managed to improve their production functions substantially in a relatively short period of time. If this argument is correct, these countries should have experienced unusually rapid growth in total factor productivity.

One study shed light on this issue by examining in detail the data from these four countries. The study found that their exceptional growth can be traced to large increases in measured factor inputs: increases in labor-force participation, increases in the capital stock, and increases in educational attainment. In South Korea, for example, the investment-GDP ratio rose from about 5 percent in the 1950s to about 30 percent in the 1980s; the percentage of the working population with at least a high-school education went from 26 percent in 1966 to 75 percent in 1991.

Once we account for growth in labor, capital, and human capital, little of the growth in output is left to explain. None of these four countries experienced unusually rapid growth in total factor productivity. Indeed, the average growth in total factor productivity in the East Asian Tigers was almost exactly the same as in the United States. Thus, although these countries' rapid growth has been truly impressive, it is easy to explain using the tools of basic growth theory.¹⁹ ■


The Solow Residual in the Short Run

When Robert Solow introduced his famous residual, his aim was to shed light on the forces that determine technological progress and economic growth in the long run. But economist Edward Prescott has looked at the Solow residual as a measure of technological change over shorter periods of time. He concludes that fluctuations in technology are a major source of short-run changes in economic activity.

Figure 8-2 shows the Solow residual and the growth in output using annual data for the United States during the period 1970 to 2007. Notice that the Solow residual fluctuates substantially. If Prescott's interpretation is correct, then we can draw conclusions from these short-run fluctuations, such as that technology worsened in 1982 and improved in 1984. Notice also that the Solow residual moves closely with output: in years when output falls, technology tends to worsen. In Prescott's view, this fact implies that recessions are driven by adverse shocks to technology. The hypothesis that technological shocks are the

¹⁹ Alwyn Young, "The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience," *Quarterly Journal of Economics* 101 (August 1995): 641-680.

FIGURE 8-2



Growth in Output and the Solow Residual The Solow residual, which some economists interpret as a measure of technology shocks, fluctuates with the economy's output of goods and services.

Source: U.S. Department of Commerce, U.S. Department of Labor, and author's calculations.

driving force behind short-run economic fluctuations, and the complementary hypothesis that monetary policy has no role in explaining these fluctuations, is the foundation for an approach called *real-business-cycle theory*.

Prescott's interpretation of these data is controversial, however. Many economists believe that the Solow residual does not accurately represent changes in technology over short periods of time. The standard explanation of the cyclical behavior of the Solow residual is that it results from two measurement problems.

First, during recessions, firms may continue to employ workers they do not need so that they will have these workers on hand when the economy recovers. This phenomenon, called *labor hoarding*, means that labor input is overestimated in recessions, because the hoarded workers are probably not working as hard as usual. As a result, the Solow residual is more cyclical than the available production technology. In a recession, productivity as measured by the Solow residual falls even if technology has not changed simply because hoarded workers are sitting around waiting for the recession to end.

Second, when demand is low, firms may produce things that are not easily measured. In recessions, workers may clean the factory, organize the inventory,

get some training, and do other useful tasks that standard measures of output fail to include. If so, then output is underestimated in recessions, which would also make the measured Solow residual cyclical for reasons other than technology.

Thus, economists can interpret the cyclical behavior of the Solow residual in different ways. Some economists point to the low productivity in recessions as evidence for adverse technology shocks. Others believe that measured productivity is low in recessions because workers are not working as hard as usual and because more of their output is not measured. Unfortunately, there is no clear evidence on the importance of labor hoarding and the cyclical mismeasurement of output. Therefore, different interpretations of Figure 8-2 persist.²⁰

MORE PROBLEMS AND APPLICATIONS

- In the economy of Solovia, the owners of capital get two-thirds of national income, and the workers receive one-third.
 - The men of Solovia stay at home performing household chores, while the women work in factories. If some of the men started working outside the home so that the labor force increased by 5 percent, what would happen to the measured output of the economy? Does labor productivity—defined as output per worker—increase, decrease, or stay the same? Does total factor productivity increase, decrease, or stay the same?
 - In year 1, the capital stock was 6, the labor input was 3, and output was 12. In year 2, the capital stock was 7, the labor input was 4, and output was 14. What happened to total factor productivity between the two years?
- Labor productivity is defined as Y/L , the amount of output divided by the amount of labor input. Start with the growth-accounting equation and show that the growth in labor pro-

ductivity depends on growth in total factor productivity and growth in the capital-labor ratio. In particular, show that

$$\frac{\Delta(Y/L)}{Y/L} = \frac{\Delta A}{A} + \alpha \frac{\Delta(K/L)}{K/L}.$$

Hint: You may find the following mathematical trick helpful. If $z = wx$, then the growth rate of z is approximately the growth rate of w plus the growth rate of x . That is,

$$\Delta z/z \approx \Delta w/w + \Delta x/x.$$

- Suppose an economy described by the Solow model is in a steady state with population growth n of 1.8 percent per year and technological progress g of 1.8 percent per year. Total output and total capital grow at 3.6 percent per year. Suppose further that the capital share of output is $1/3$. If you used the growth-accounting equation to divide output growth into three sources—capital, labor, and total factor productivity—how much would you attribute to each source? Compare your results to the figures we found for the United States in Table 8-3.

²⁰ To read more about this topic, see Edward C. Prescott, “Theory Ahead of Business Cycle Measurement,” and Lawrence H. Summers, “Some Skeptical Observations on Real Business Cycle Theory,” both in *Quarterly Review*, Federal Reserve Bank of Minneapolis (Fall 1986); N. Gregory Mankiw, “Real Business Cycles: A New Keynesian Perspective,” *Journal of Economic Perspectives* 3 (Summer 1989): 79–90; Bennett T. McCallum, “Real Business Cycle Models,” in R. Barro, ed., *Modern Business Cycle Theory* (Cambridge, MA: Harvard University Press, 1989), 16–50; and Charles I. Plosser, “Understanding Real Business Cycles,” *Journal of Economic Perspectives* 3 (Summer 1989): 51–77.