the capital stock and labor supply and ultimately determine economic output. New saving generates additional investment, which adds to the capital stock and allows for increased output. New workers add further to the economy's capacity to increase production.

One way these five equations can be simplified slightly is to combine equations 4–2, 4–3, and 4–4. The aggregate level of saving (in equation 4–2) determines the level of investment in equation 4–3, which (together with depreciation) determines changes in the capital stock in equation 4–4. Combining these three equations gives us

$$\Delta K = sY - dK \qquad [4\text{–}6]$$

This equation states that the change in the capital stock (ΔK) is equal to saving (sY) minus depreciation (dK). This expression allows us to calculate the change in the capital stock and enter the new value directly into the aggregate production function in equation 4–1.

THE HARROD-DOMAR GROWTH MODEL

As we have stressed, the aggregate production function (shown earlier as equation 4–1) is at the heart of every model of economic growth. This function can take many different forms, depending on what we believe is the true relationship between the factors of production (K and L) and aggregate output. This relationship depends on (among other things) the mix of economic activities (for example, agriculture, heavy industry, light labor-intensive manufacturing, high-technology processes, services), the level of technology, and other factors. Indeed, much of the theoretical debate in the academic literature on economic growth is about how to best represent the aggregate production process.

THE FIXED-COEFFICIENT PRODUCTION FUNCTION

One special type of a simple production function is shown in Figure 4–1. Output in this figure is represented by **isoquants**, which are combinations of the inputs (labor and capital in this case) that produce equal amounts of output. For example, on the first (innermost) isoquant, it takes capital (plant and equipment) of \$10 million and 100 workers to produce 100,000 keyboards per year (point a). Alternatively, on the second isoquant, \$20 million of capital and 200 workers can produce 200,000 keyboards (point b). Only two isoquants are shown in this diagram, but a nearly infinite number of isoquants are possible, each for a different level of output.

The L-shape of the isoquants is characteristic of a particular type of production function known as **fixed-coefficient production functions**. These production

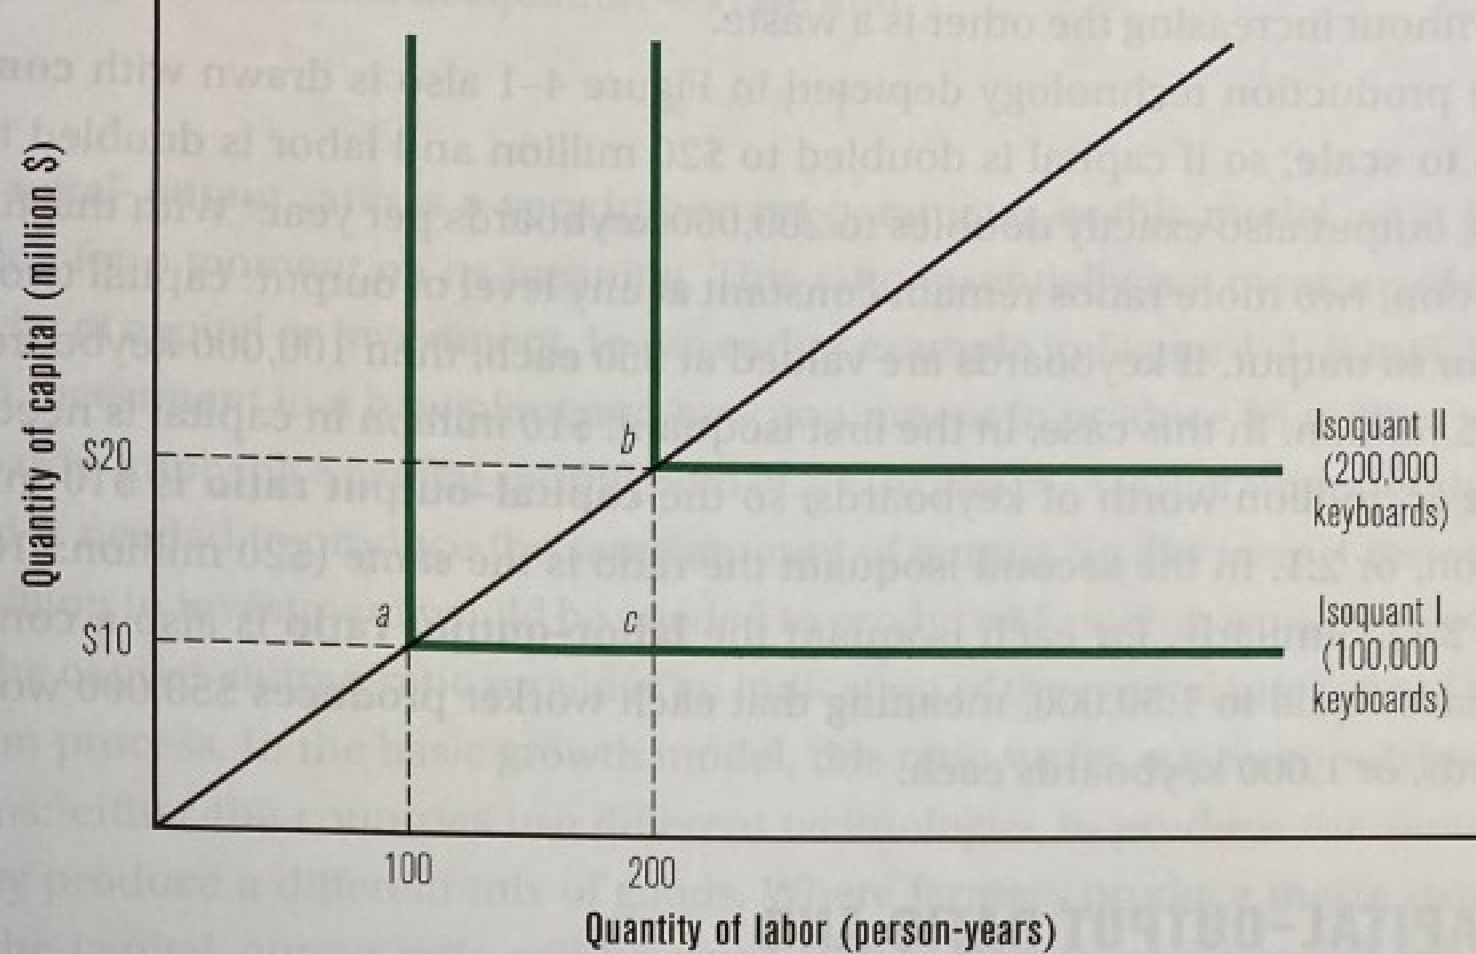

FIGURE 4–1 Isoquants for a Fixed-Coefficient Production Technology

With constant returns to scale, the isoquants will be L-shaped and the production function will be the straight line through their minimum-combination points.

functions are based on the assumption that capital and labor need to be used in a fixed proportion to each other to produce different levels of output. In Figure 4–1, for the first isoquant, the **capital–labor ratio** is 10 million:100, or 100,000:1. In other words, $100,000 in capital must be matched with one worker to produce the given output. For the second isoquant, the ratio is the same: $20 million:200, or 100,000:1. This constant capital–labor ratio is represented in Figure 4–1 by the slope of the ray from the origin through the vertices (*a* and *b*) of the isoquants. These vertices represent the least cost and hence most efficient mix of capital and labor to produce a given quantity of output. In the case of fixed coefficients, this mix is the same for every quantity of output.

With this kind of production function, if more workers are added *without* investing in more capital, output does *not* rise. Look again at the first isoquant, starting at the elbow (with 100 workers and $10 million in capital). If the firm adds more workers (say, increasing to 200 workers) without adding new machines, it moves horizontally to the right along the first isoquant to point *c*. But at this point, or at any other point on this isoquant, the firm still produces just 100,000 keyboards. In this kind of production function, new workers need more machines to increase output. Adding new workers without machines results in idle workers, with no increase in output. Similarly, more machinery without additional workers results in underused machines. On each isoquant, the most efficient production point is at the elbow, where the minimum amounts of capital and labor are used. To use any more of either factor without increasing the other is a waste.

The production technology depicted in Figure 4–1 also is drawn with **constant returns to scale**, so if capital is doubled to \$20 million and labor is doubled to 200 workers, output also exactly doubles to 200,000 keyboards per year.[6] With this further assumption, two more ratios remain constant at any level of output: capital to output and labor to output. If keyboards are valued at \$50 each, then 100,000 keyboards are worth \$5 million. In this case, in the first isoquant, \$10 million in capital is needed to produce \$5 million worth of keyboards, so the **capital–output ratio** is \$10 million:\$5 million, or 2:1. In the second isoquant the ratio is the same (\$20 million:\$10 million, or 2:1). Similarly, for each isoquant the **labor–output ratio** is also a constant, in this case equal to 1:50,000, meaning that each worker produces \$50,000 worth of keyboards, or 1,000 keyboards each.

THE CAPITAL–OUTPUT RATIO AND THE HARROD-DOMAR FRAMEWORK

The fixed-coefficient, constant-returns-to-scale production function is the centerpiece of a well-known early model of economic growth that was developed independently during the 1940s by economists Roy Harrod of England and Evsey Domar of MIT, primarily to explain the relationship between growth and unemployment in advanced capitalist societies.[7] It ultimately focuses attention on the role of capital accumulation in the growth process. The **Harrod-Domar model** has been used extensively (perhaps even overused) in developing countries to examine the relationship between growth and capital requirements. The model is based on the real-world observation that some labor is unemployed and proceeds on the basis that capital is the binding constraint on production and growth. In the model, the production function has a very precise form, in which output is assumed to be a *linear* function of capital (and only capital). As usual, the model begins by specifying the level of output, which we later modify to explore changes in output, or economic growth. The production function is specified as follows:

$$Y = 1/v \times K \quad or \quad Y = K/v \qquad [4\text{–}7]$$

where v is a constant. In this equation, the capital stock is multiplied by the fixed number $1/v$ to calculate aggregate production. If $v = 3$ and a firm has \$30 million in capital, its annual output would be \$10 million. It is difficult to imagine a simpler

[6] In a constant-returns-to-scale production function, if we multiply both capital and labor by any number, w, output multiplies by the same number. In other words, the production function has the following property: $wY = F(wK, wL)$.

[7] Roy F. Harrod, "An Essay in Dynamic Theory," *Economic Journal* (1939), 14–33; Evsey Domar, "Capital Expansion, Rate of Growth, and Employment," *Econometrica* (1946), 137–47; and Domar, "Expansion and Employment," *American Economic Review* 37 (1947), 34–55.

production function. The constant v turns out to be the capital–output ratio because, by rearranging the terms in equation 4–7, we find

$$v = K/Y \tag{4–8}$$

The capital–output ratio is a very important parameter in this model, so it is worth dwelling for a moment on its meaning. This ratio essentially is a measure of the productivity of capital or investment. In the earlier example in Figure 4–1, it took $10 million in investment in a new plant and new equipment to produce $5 million worth of keyboards, implying a capital–output ratio of 2:1 (or just 2). A larger v implies that more capital is needed to produce the same amount of output. So, if v were 4 instead, then $20 million in investment would be needed to produce $5 million worth of keyboards.

The capital–output ratio provides an indication of the capital intensity of the production process. In the basic growth model, this ratio varies across countries for two reasons: either the countries use different technologies to produce the same goods or they produce a different mix of goods. Where farmers produce maize using tractors, the capital–output ratio will be much higher than in countries where farmers rely on a large number of workers using hoes and other hand tools. In countries that produce a larger share of **capital-intensive products** (that is, those that require relatively more machinery, such as automobiles, petrochemicals, and steel), v is higher than in countries producing more **labor-intensive products** (such as textiles, basic agriculture, and foot wear). In practice, as economists move from the v of the model to actually measuring it in the real world, the observed capital–output ratio can also vary for a third reason: differences in efficiency. A larger measured v can indicate less-efficient production when capital is not being used as productively as possible. A factory with lots of idle machinery and poorly organized production processes has a higher capital–output ratio than a more-efficiently managed factory.

Economists often calculate the **incremental capital–output ratio (ICOR)** to determine the impact on output of additional (or incremental) capital. The ICOR measures the productivity of additional capital, whereas the (average) capital–output ratio refers to the relationship between a country's total stock of capital and its total national product. In the Harrod-Domar model, because the capital–output ratio is assumed to remain constant, the average capital–output ratio is equal to the incremental capital–output ratio, so the ICOR $= v$.

So far, we have been discussing total output, not growth in output. The production function in equation 4–7 easily can be converted to relate *changes* in output to *changes* in the capital stock:

$$\Delta Y = \Delta K/v \tag{4–9}$$

The growth rate of output, g, is simply the increment in output divided by the total amount of output, $\Delta Y/Y$. If we divide both sides of equation 4–9 by Y, then

$$g = \Delta Y/Y = \Delta K/Yv \tag{4–10}$$

Finally, from equation 4–6, we know that the change in the capital stock ΔK is equal to saving minus the depreciation of capital ($\Delta K = sY - dK$). Substituting the right-hand side of equation 4–6 into the term for ΔK in equation 4–10 and simplifying[8] leads to the basic Harrod-Domar relationship for an economy:

$$g = (s/v) - d \qquad [4\text{–}11]$$

Underlying this equation is the view that capital created by investment is the main determinant of growth in output and that saving makes investment possible.[9] It rivets attention on two keys to the growth process: saving (s) and the productivity of capital (v). The message from this model is clear: Save more and make productive investments, and your economy will grow.

Economic analysts can use this framework either to predict growth or to calculate the amount of saving required to achieve a target growth rate. The first step is to try to estimate the incremental capital–output ratio (v) and depreciation rate (d). With a given saving rate, predicting the growth rate is straightforward. If the saving (or investment) rate is 24 percent, the ICOR is 3, and the depreciation rate is 5 percent, then the economy can be expected to grow by 3 percent (because $0.24/3 - 0.05 = 0.03$).

How does this model work in practice? Consider Indonesia, which from 2002 to 2007 had an investment rate of about 30 percent and recorded a GDP growth just under 5.5 percent per year. Assuming a depreciation rate of 5 percent, the implied ICOR was approximately $v = 2.86$.[10] Would these figures have helped the Indonesian government predict the 2007–08 growth rate? In 2008, the investment rate was 29 percent, so the Harrod-Domar model would have predicted growth of 5.1 percent ($g = 0.29/2.86 - 0.05$). The actual growth rate in 2007–08 was 6.0 percent, within sight of the prediction but not highly accurate.

STRENGTHS AND WEAKNESSES OF THE HARROD-DOMAR FRAMEWORK

The basic strength of the Harrod-Domar model is its simplicity. The data requirements are small, and the equation is easy to use and to estimate. And, as we saw with the example of Indonesia, the model can be somewhat accurate from one year to the next. Generally speaking, in the absence of severe economic shocks (such as a drought, a financial crisis, or large changes in export or import prices), the model can do a reasonable job of estimating expected growth rates in most countries over very short periods of time (a few years). Another strength is its focus on the key role of saving. As discussed

[8]Substituting equation 4–6 into 4–10 leads to $g = (sY - dK)/Y \times 1/v$, which can be simplified to $g = (s - d \times K/Y) \times 1/v$. Since $K/Y = v$, we have $g = (s - dv) \times 1/v$, which leads to $g = s/v - d$.

[9]For an important early contribution to the discussion of the importance of capital accumulation to the growth process, see Joan Robinson, *The Accumulation of Capital* (London: Macmillan, 1956).

[10]Because $g = s/v - d$, then $v = s/(g + d)$. For Indonesia between 2002 and 2007, $v = 0.30/(0.055 + 0.05) = 2.86$.

in Chapter 3, individual decisions about how much income to save and consume are central to the growth process. People prefer to consume sooner rather than later, but the more that is consumed, the less that can be saved to finance investment. The Harrod-Domar model makes it clear that saving is crucial for income to grow over time.

The model, however, has some major weaknesses. One follows directly from the strong focus on saving. Although saving is necessary for growth, the simple form of the model implies that it is also sufficient, which it is not. As pointed out in Chapter 3, the investments financed by saving actually have to pay off with higher income in the future, and not all investments do so. Poor investment decisions, changing government policies, volatile world prices, or simply bad luck can alter the impact of new investment on output and growth. Sustained growth depends both on generating new investment and ensuring that investments are productive over time. In this vein, the allocation of resources across different sectors and firms can be an important determinant of output and growth. Because (for simplicity) the Harrod-Domar assumes only one sector, it leaves out these important allocation issues.

Perhaps the most important limitations in the model stem from the rigid assumptions of fixed capital-to-labor, capital-to-output, and labor-to-output ratios, which imply very little flexibility in the economy over time. To keep these ratios constant, capital, labor, and output must all grow at exactly the same rate, which is highly unlikely to happen in real economies. To see why these growth rates must all be the same, consider the growth rate of capital. If the capital stock grew any faster or slower than output at rate g, the capital–output ratio would change. Thus the capital stock must grow at g to keep the capital–output ratio constant over time. With respect to labor, in our original five-equation model, we stipulated (in equation 4–5) that the labor force would grow at exactly the same pace as the population at rate n. Therefore, the only way that the capital stock and the labor force can grow at the same rate is if n happens to be equal to g. This happens only when $n = g = s/v - d$, and there is no particular reason to believe the population will grow at that rate.

In this model, the economy remains in equilibrium with full employment of the labor force and the capital stock *only* under the very special circumstances that labor, capital, and output all *grow* at the rate g. On the one hand, if n is larger than g, the labor force grows faster than the capital stock. In essence, the saving rate is not high enough to support investment in new machinery sufficient to employ all new workers. A growing number of workers do not have jobs and unemployment rises indefinitely. On the other hand, if g (or $s/v - d$) is larger than n, the capital stock grows faster than the workforce. There are not enough workers for all the available machines, and capital becomes idle. The actual growth rate of the economy no longer is g, as the model stipulates, but slows to n, with output constrained by the number of available workers.

So, unless $s/v - d$ (or g) is exactly equal to n, either labor or capital is not fully employed and the economy is not in a stable equilibrium. This characteristic of the Harrod-Domar model has come to be known as the *knife-edge* problem. As long as

$g = n$, the economy remains in equilibrium, but as soon as either the capital stock or the labor force grows faster than the other, the economy falls off the edge with continuously growing unemployment of either capital or labor.

The rigid assumptions of fixed capital–output, labor–output, and capital–labor ratios may be reasonably accurate for short periods of time or in very special circumstances but almost always are inaccurate over time as an economy evolves and develops. Each of these varies among countries and, for a single country, over time. Consider the incremental capital–output ratio. The productivity of capital can change in response to policy changes, which in turn affects v. Moreover, the capital intensity of the production process can and usually does change over time. A poor country with a low saving rate and surplus labor (unemployed and underemployed workers) can achieve higher growth rates by utilizing as much labor as possible and thus relatively less capital. For example, a country relying heavily on labor-intensive agricultural production will record a low v. As economies grow and per capita income rises, the labor surplus diminishes and economies shift gradually toward more capital-intensive production. As a result, the ICOR shifts upward. Thus a higher v may not necessarily imply inefficiency or slower growth. ICORs can also shift through market mechanisms, as prices of labor and capital change in response to changes in supplies. As growth takes place, saving tends to become relatively more abundant and hence the price of capital falls while employment and wages rise. Therefore, all producers increasingly economize on labor and use more capital and the ICOR tends to rise.

Consider again the example of Indonesia. The ICOR changed from approximately 2.4 during the 1980s, to 4.1 during the 1990s, to 3.6 during 2000–09, reflecting a trend toward more capital-intensive production processes. To continue to use the 1980s ICOR in 2009 would have been very misleading and betrayed a significant misunderstanding of the growth process. The structure of the economy had changed substantially during that time period and, with it, the ICOR. Thailand provides a similar example, as described in Box 4–1.

As a result of these rigidities, the Harrod-Domar framework tends to become increasingly inaccurate over longer periods of time as the actual ICOR changes and, with it, the capital–labor ratio. In a world with fixed-coefficient production functions, little room is left for a factory manager to increase output by hiring one more worker without buying a machine to go with the worker or to purchase more machines for the current workforce to use. The fixed-proportion production function does not allow for any substitution between capital and labor in the production process. In the real world, of course, at least some substitution between labor and capital is possible in most production processes. As we see in the next section, adding this feature to the model allows for a much richer exploration of the growth process.

A final weakness of the Harrod-Domar model is the absence of any role for productivity growth—the ability to produce increasing quantities of output per unit of input. In Figure 4–1, increased factor productivity can be represented by an inward shift of each isoquant toward the origin, implying that less labor and capital would

BOX 4-1 ECONOMIC GROWTH IN THAILAND

In the 1960s, Thailand's agrarian economy depended heavily on rice, maize, rubber, and other agricultural products. About three-quarters of the Thai population derived its income from agricultural activities. Gross domestic product (GDP) per capita in 1960 (measured in 2005 international dollars) was around $1,200—less than one-tenth the average income in the United States. Life expectancy was 53 years and the infant mortality rate was 103 per 1,000 births. Few observers expected Thailand to develop rapidly.

However, since the mid-1960s, the Thai economy has grown rapidly (if not always steadily), benefiting from relatively sound economic management and a favorable external environment. The government regularly achieved surpluses on the current account of its budget and used these funds (plus modest inflows of foreign assistance) to finance investments in rural roads, irrigation, power, telecommunications, and other basic infrastructure. At least until the mid-1990s, the government's fiscal, monetary, and exchange rate policies kept the macroeconomy relatively stable with fairly low inflation, despite the turbulent period of world oil price shocks in the 1970s and 1980s. Beginning in the 1970s, the government began to remove trade restrictions and promote the production of labor-intensive manufactured exports. These products found a ready market in the booming Japanese economy of the 1980s and provided a growing number of jobs for Thai workers.

Thailand's ability to make investments and deepen its capital stock depended on its capacity to save. The country's saving rate averaged about 20 percent in the 1960s, already high for developing countries, and increased steadily over time to an average of 35 percent in the 1990s, falling to about 32 percent during the 2000s. These high saving rates, combined with relatively prudent economic policies, supported very rapid economic growth and development.

Thailand's development experience has been far from completely smooth, however. In mid-1997, a major financial crisis erupted. Huge short-term offshore borrowing combined with a fixed exchange rate and weak financial institutions led to a collapse of a real estate bubble, rapid capital flight, a substantial depreciation of the Thai baht, and a deep recession (see Chapter 13). In some ways, Thailand had become the victim of its own success, with its rapid growth attracting significant numbers of investors looking to gain quick profits, who rapidly fled once the bubble began to collapse. After two years of negative growth (with GDP falling 10 percent in 1998), the economy began to recover and growth rebounded to 3.9 percent between 1999 and 2007.

Over the longer period between 1960 and 2007, per capita growth averaged 4.5 percent, so that the average income in Thailand is now more than eight times higher than it was in 1960. By 2009, life expectancy grew to 69 years, infant mortality fell to 12 per thousand, and adult literacy reached 93 percent. During this period, the structure of the economy changed significantly. By 2009, manufacturing accounted for 34 percent of GDP, up from just 14 percent in 1965, while the share of agricultural production dropped commensurately. The composition of exports shifted away from rice, maize, and other agricultural commodities toward labor-intensive manufactured products, which now account for a large majority of all exports. As the Harrod-Domar model predicts, Thailand's high saving rate and resulting capital accumulation was accompanied by a dramatic increase in output (and income) per capita. Contrary to the Harrod-Domar model, however, the ICOR did not remain constant. As the stock of capital grew and the economy shifted toward more capital-intensive production techniques, the ICOR increased from 2.6 in the 1970s to nearly 5 by the early 2000s. The rising ICOR indicated that, as the Thai economy expanded and the level of capital per worker increased, an ever-larger increment of new capital was required to bring about a given increase in total output.

be needed to produce the same amount of output. The simplest way to capture this in the Harrod-Domar framework is to introduce a smaller ICOR, but of course, this would contradict the idea of a constant ICOR.

Despite these weaknesses, the Harrod-Domar model is still used to a surprisingly wide extent. Economist William Easterly documented how the World Bank and other institutions use the model to calculate "financing gaps" between the amount of available saving and the amount of investment supposedly needed to achieve a target growth rate.[11] He shows how simplistic and sometimes careless use of the model can lead to weak analysis and faulty conclusions. In essence, analysts enamored by the simplicity of the model tend to overlook its shortcomings when applying it to the real world.

The Harrod-Domar model provides some useful insights but does not take us very far. The fixed-coefficient assumption provides the model with very little flexibility and does not capture the ability of real world firms to change the mix of inputs in the production process. The model can be reasonably accurate from one year to the next (in the absence of shocks), and it rightly focuses attention on the importance

[11]See William Easterly, "Aid for Investment," *The Elusive Quest for Growth* (Cambridge: MIT Press, 2001), chap. 2; and Easterly, "The Ghost of the Financing Gap: Testing the Growth Model of the International Financial Institutions," *Journal of Development Economics* 60, no. 2 (December 1999), 423–38.

of saving. But it is quite inaccurate for most countries over longer periods of time and implies that saving is sufficient for growth, although it is not. Indeed, in the late 1950s, Domar expressed strong doubts about his own model, pointing out that it was originally designed to explore employment issues in advanced economies rather than growth per se and was too rigid to be useful for explaining long-term growth.[12] Instead, he endorsed the new growth model of Robert Solow, to which we now turn our attention.

THE SOLOW (NEOCLASSICAL) GROWTH MODEL

THE NEOCLASSICAL PRODUCTION FUNCTION

In 1956, MIT-economist Robert Solow introduced a new model of economic growth that was a big step forward from the Harrod-Domar framework.[13] Solow recognized the problems that arose from the rigid production function in the Harrod-Domar model. Solow's answer was to drop the fixed-coefficients production function and replace it with a **neoclassical production function** that allows for more flexibility and substitution between the factors of production. In the Solow model, the capital–output and capital–labor ratios no longer are fixed but vary, depending on the relative endowments of capital and labor in the economy and the production process. Like the Harrod-Domar model, the Solow model was developed to analyze industrialized economies, but it has been used extensively to explore economic growth in all countries around the world, including developing countries. The Solow model has been enormously influential and remains at the core of most theories of economic growth in developing countries.

The isoquants that underlie the neoclassical production function are shown in Figure 4–2. Note that the isoquants are curved rather than L-shaped as in the fixed-coefficient model. In this figure, at point *a*, \$10 million of capital and 100 workers combine to produce 100,000 keyboards, which would be valued at \$5 million (because, as stated earlier, keyboards are priced at \$50 each). Starting from this point, output could be expanded in any of three ways. If the firm's managers decided to expand at constant factor proportions and move to point *b* on isoquant II to produce 200,000 keyboards, the situation would be identical to the fixed propor-

[12] Evsey Domar, *Essays in the Theory of Economic Growth* (Oxford: Oxford University Press, 1957).

[13] The two classic references of Solow's work are his "A Contribution to the Theory of Economic Growth" and "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics* 39 (August 1957), 312–20. For an excellent and very thorough undergraduate exposition of the Solow and other models of economic growth, see Charles I. Jones, *Introduction to Economic Growth* (New York: W. W. Norton and Company, 2001). In 1987, Solow was awarded the Nobel Prize in economics, primarily for his work on growth theory.