Fletcher School, Tufts University

17. Capital Accumulation, Technical Progress and Economic Growth

E212 Macroeconomics

Prof. George Alogoskoufis

From the Classical to the Neoclassical Theory of Growth

The analysis of long run economic growth, along with the distribution of income, was an important concern of the so called *classical economists* such as Smith, Malthus, Ricardo and Mill.

The classical economists sought to explain differences in the wealth of nations and economic growth in terms of population growth, the accumulation of capital and the increase in the efficiency of production, as determined by the division of labor and technical progress. These factors, interacted with land, a factor of production which was assumed to be in fixed supply. The classical theory of economic growth, focusing on the production and distribution of income, was essentially macroeconomic in nature.

Modern macroeconomics use a one sector neoclassical model, developed originally by Robert Solow, from the Massachusetts Institute of Technology (MIT) in the mid-1950s.

This model highlights the role of substitutability between capital and labor in the production process, the accumulation of physical and human capital, and the role of savings and investment, population growth and technical progress. It assumes that the technology of production, the savings rate, the rate of population growth and the rate of technical progress are exogenously given, and it studies how they affect living standards and economic growth.

Production Technology: The Aggregate Production Function

The starting point for any theory of growth must be the technological process which transforms the inputs of production into aggregate output. This process is summarized by an aggregate production function of the form,

$$Y=F(K,N)$$

Y is aggregate output, K is the stock of capital (including land) and N is employment of labor. The function F, the *production function*, determines how much output is produced for given quantities of capital and labor.

The aggregate production function we introduced to study the determination of output in the short run and the medium run took a particularly simple form. Output was simply proportional to the amount of labor used by firms—more specifically, proportional to the number of workers employed by firms.

$$Y=AN$$

where A was a constant denoting the productivity of labor. So long as our focus was on fluctuations in output and employment, the assumption was acceptable. But now that our focus has shifted to growth this assumption will no longer do: It implies that output per worker is constant, ruling out growth (or at least growth of output per worker) altogether. It is time to relax it. From now on, we will assume that there are two inputs—capital and labor—and that the relation between aggregate output and the two inputs is given by the production function F.

Some Thoughts on Technology, Institutions and the Aggregate Production Function

Assuming the existence of an aggregate production function is a dramatic simplification of reality. Surely, machines and office buildings play very different roles in production and should be treated as separate inputs. Surely, workers with Ph.D.'s are different from high-school dropouts; yet, by constructing the labor input as simply the number of workers in the economy, we treat all workers as identical.

Yet the production function has proven a useful way to think about the technological relation between aggregate inputs and output in an economy. A country with a more advanced technology will produce more output from the same quantities of capital and labor than an economy with a primitive technology.

How should we define the state of technology? Should we think of it as the list of blueprints defining both the range of products that can be produced in the economy as well as the techniques available to produce them? Or should we think of it more broadly, including not only the list of blueprints, but also the way the economy is organized—from the internal organization of firms, to the system of laws and the quality of their enforcement, to the political system, and so on? This broader way of thinking about the technology of production has much to commend it, but we shall postpone all further discussion for a while, taking the state of technology in any country as exogenously given.

The Neoclassical Production Function

A number of important assumptions are made in order to describe the properties of the aggregate production function.

1. Output increases when any of the factors of production increases. Thus, higher capital for given labor results in higher output, and higher labor for given capital also results in higher output. Thus, the *marginal product* of both factors of production is positive.

2. The extra output produced by an extra unit of either capital or labor is declining as capital or labor increases. Thus, the *marginal product of both factors of production is positive but declining* in the utilization of each particular factor.

3. Raising both factors of production by the same proportion, will result in a rise in output by the same proportion. For example, if both factors of production are doubled, output will also double. This assumption is known as *constant returns to scale*.

Thus, the aggregate production function is characterized by constant returns to scale and positive but declining marginal products for each factor of production. A production function with these properties is known as the neoclassical production function.

The Neoclassical Production Function in Algebraic Form

With the assumptions we have made, we can write the neoclassical production function as a function expressed in terms of output per worker.

Note that constant returns to scale imply that if we multiply both factors of production by λ then output will also be multiplied by λ.

$$\lambda Y = F(\lambda K, \lambda N)$$

If we set $\lambda = 1/N$, then if follows that,

$$Y/N = F(K/N, 1)$$

Thus, the neoclassical production function implies that output per worker will be a positive but declining function of capital per worker.

As the capital stock increases relative to labor, output also increases. However, it increases at a declining rate because of the positive but declining marginal product of capital.

The Neoclassical Production Function in Diagrammatic Form

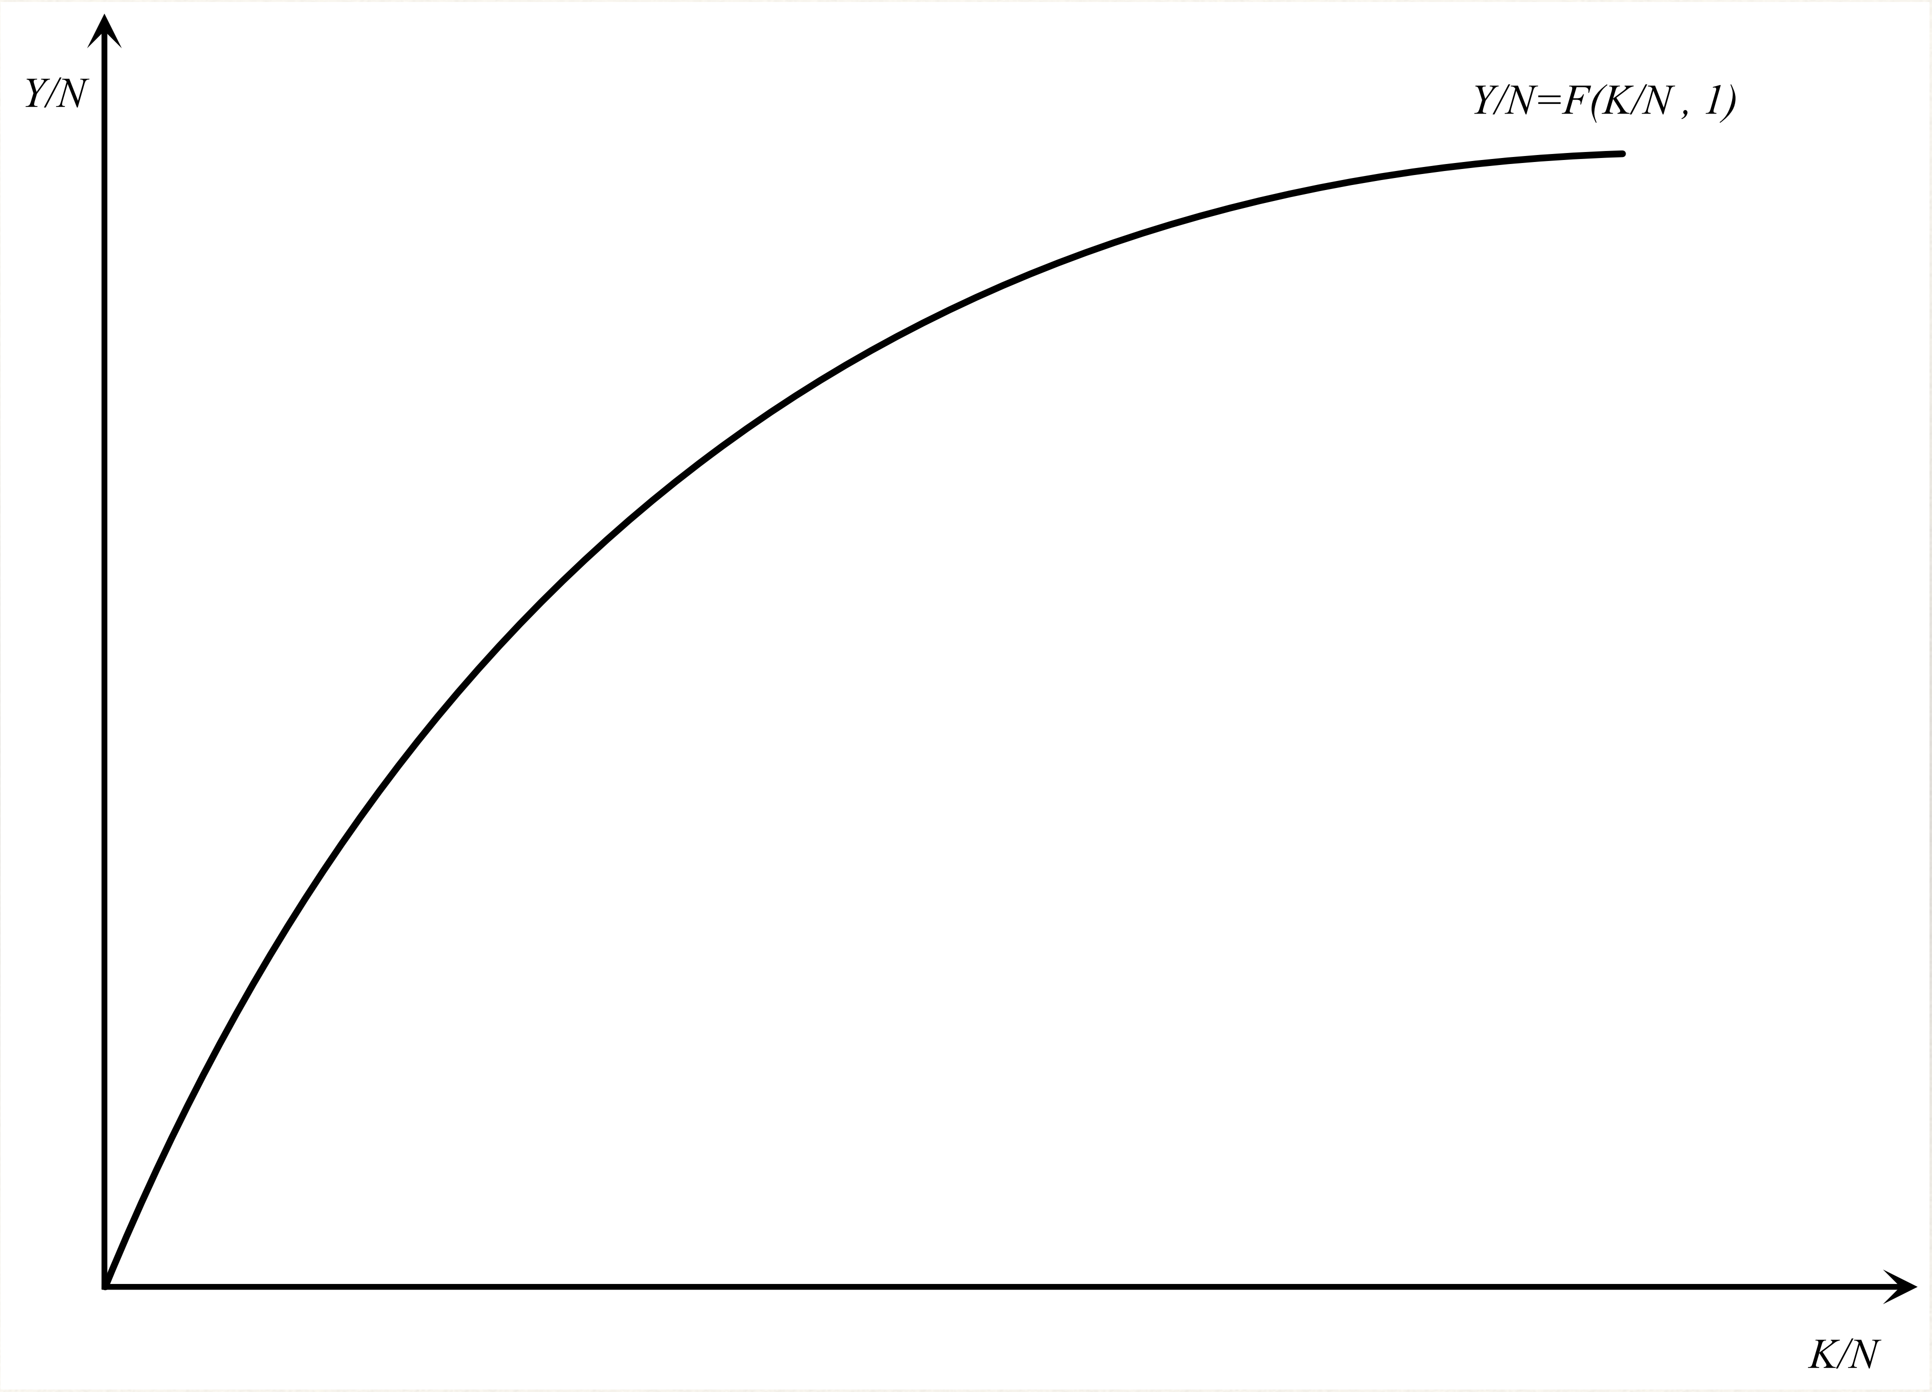

The Sources of Growth in Per Capita Output

We are now ready to return to our basic question: Where does growth come from? Why does output per worker—or output per person, if we assume the ratio of workers to the population as a whole remains constant over time—go up over time? The neoclassical production function gives a first answer:

Increases in output per worker *(Y/N)* can come from increases in capital per worker (K/N). This is the relation we just looked at in the previous figure. As (K/N) increases, that is, as we move to the right on the horizontal axis, *(Y/N)* increases.

Increases in output per worker can also come from improvements in the state of technology that shift the production function *F* upwards, and lead to more output per worker given capital per worker. This is shown in the next Figure. An improvement in the state of technology shifts the production function up, from *F(K/N , 1)* to *F'(K/N , 1)*. For a given level of capital per worker, the improvement in technology leads to an increase in output per worker. For example, for the level of capital per worker corresponding to point *A*, output per worker, increases from *A* to *A'*.

Technical Progress and Economic Growth

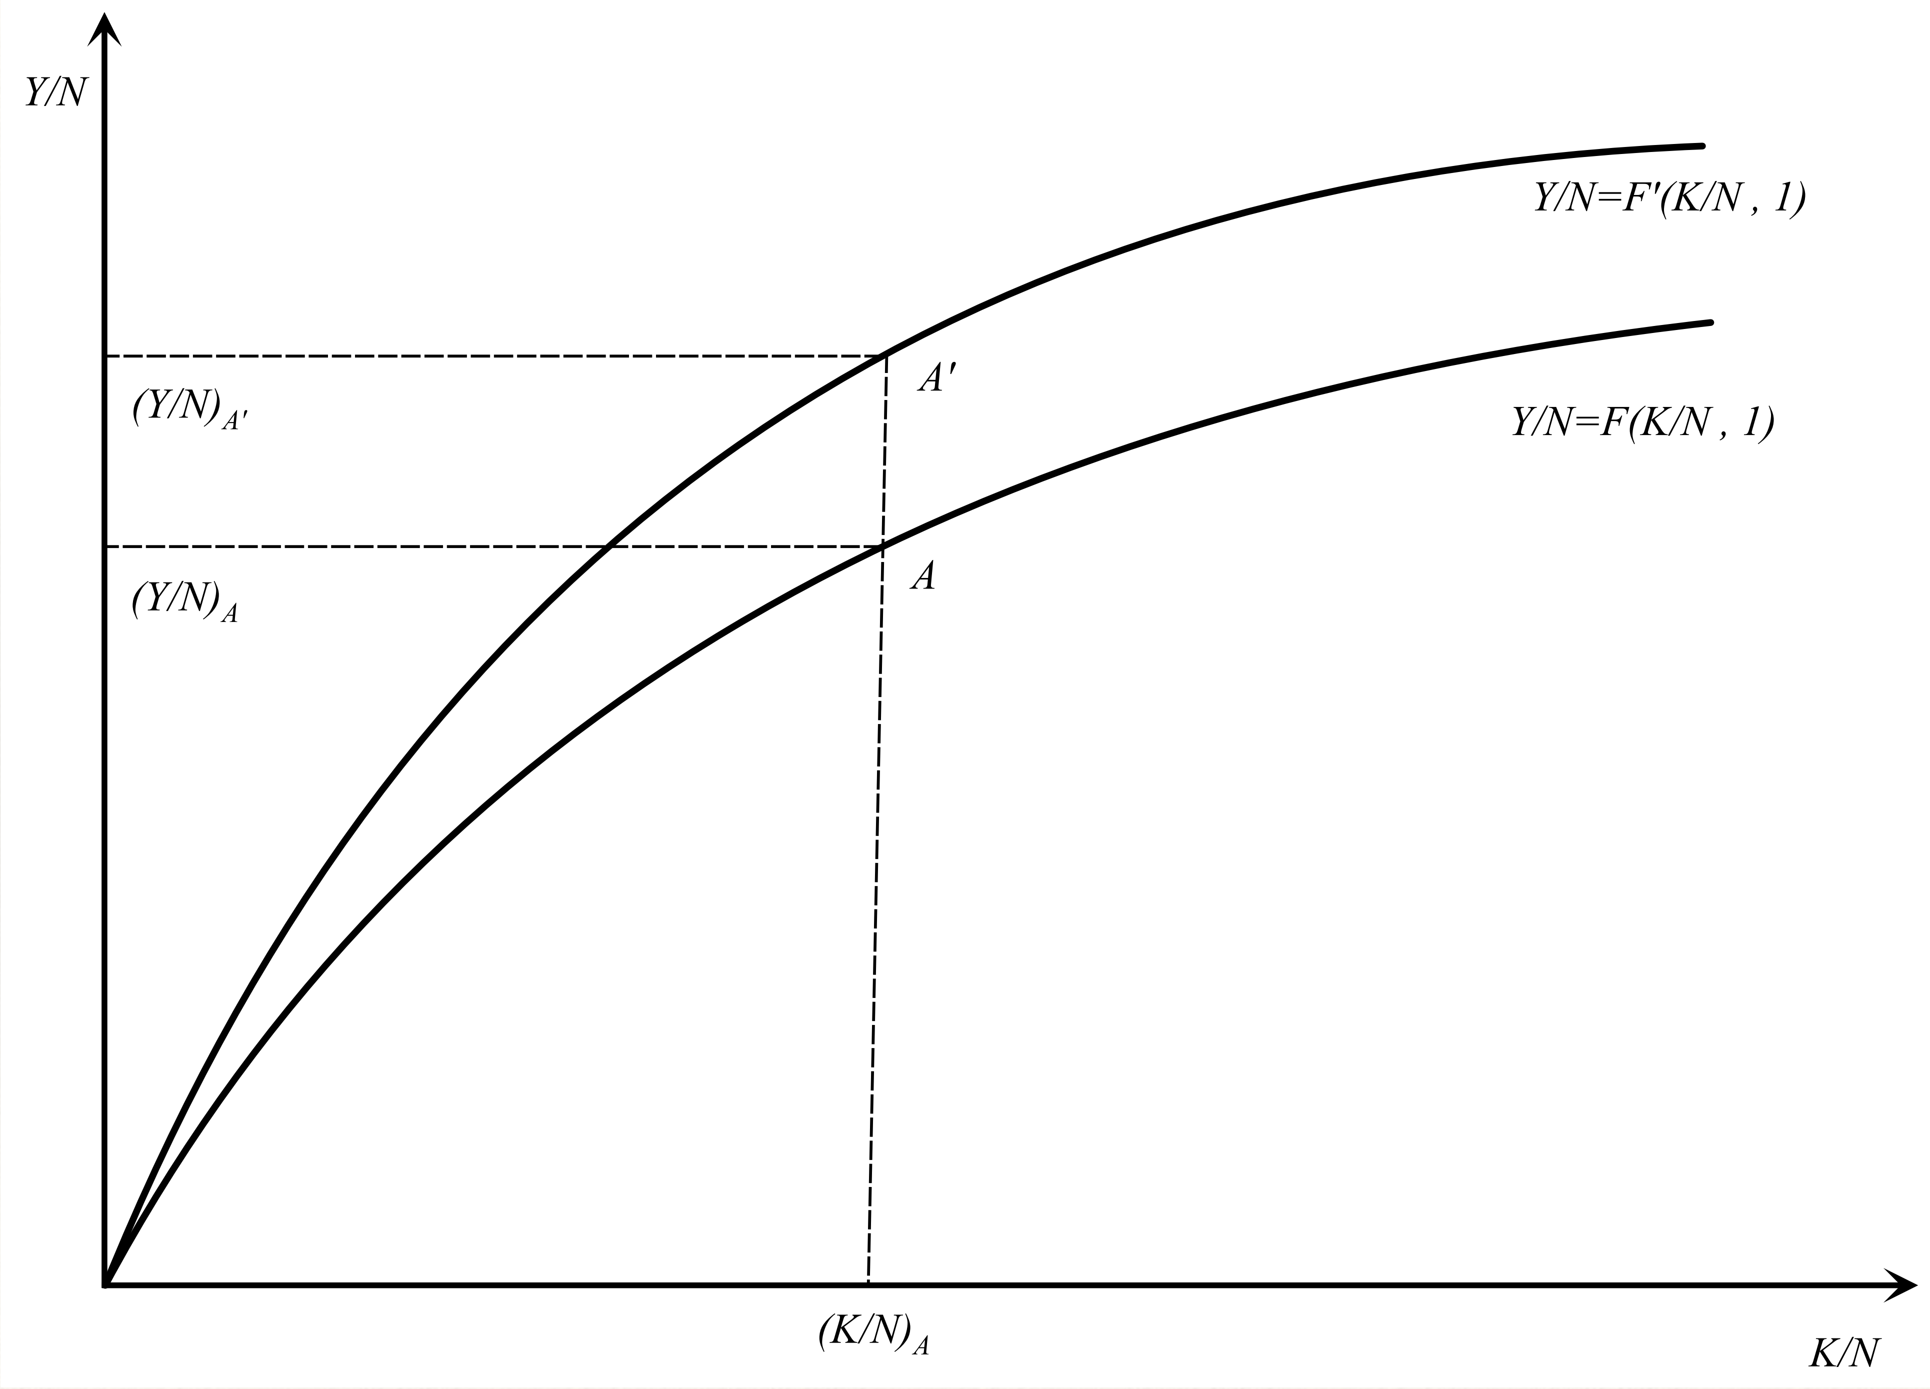

Capital Accumulation versus Technical Progress

Our analysis of the production function leads us to the conclusion that growth in per capita output could result from either capital accumulation or from technical progress—the improvement in the state of technology. We will see, however, that these two factors play very different roles in the growth process:

Capital accumulation by itself cannot sustain growth. Because of decreasing returns to capital, sustaining a steady increase in output per worker will require larger and larger increases in the level of capital per worker. At some stage, the economy will be unwilling or unable to save and invest enough to further increase capital. At that stage, output per worker will stop growing.

Does this mean that an economy's saving rate, the proportion of income that is saved and invested, is irrelevant? No. It is true that a higher saving rate cannot permanently increase the growth rate of output. But, as we shall soon see, a higher saving rate can sustain a higher level of output.

Sustained growth requires sustained technical progress. This really follows from the previous proposition: Given that the two factors that can lead to an increase in output are capital accumulation and technical progress, if capital accumulation cannot sustain growth forever, then technical progress must be the key to growth. And it is. We will soon see an economy's rate of growth of output per person is eventually determined by its rate of technical progress.

The Simple Mathematics of Cumulative Growth

Before we delve deeper in the economics of economic growth it is useful to review the simple mathematics of cumulative growth.

Assume you have a variable x which grows over time at a rate $g>0$. If we use t as an index of time, it follows that after one period we shall have that,

$$x_{t+1}=(1+g)x_t$$

After 2 periods we shall have that,

$$x_{t+2}=(1+g)x_{t+1}=(1+g)^2x_t$$

After T periods we shall have that,

$$x_{t+T}=(1+g)x_{t+T-1}=(1+g)^Tx_t$$

As time goes by, the variable grows continuously at a rate g.

How many periods will it take for the variable to double in size. This depends on the growth rate. The higher the growth rate, the fewer periods it will take for the variable to double in size.

Take two extremes. At a zero growth rate, the variable will remain constant over time. At a growth rate of 100% in will double in size after only one period. Typically growth rates are rather small, in the region between say 1% and 10%. How many periods will it take then for the variable to double at such growth rates.

The Rule of 70

For a variable x which grows over time at a rate $g>0$, we have that after T periods it will be equal to,

$$x_{t+T}=(1+g)^T x_t$$

How many years will it take for the variable to double if it grows at a rate of $g\%$ per period. To find this we must solve the equation,

$$x_{t+T}=(1+g)^T x_t=2x_t$$

This implies that,

$$(1+g)^T=2$$

To solve this equation for T, take the logarithm of both sides. It follows that,

$$T\ \ln(1+g)=\ln(2)$$

where ln denotes natural logarithms. Since for a small g, from the properties of logarithms $\ln(1+g) \simeq g$, it follows that it is easy to calculate the number of years it takes for x to double at different growth rates. It is given by,

$$T=\ln(2)/\ln(1+g) \simeq 0.693/g = 69.3/g\%$$

Hence, the number of periods required for x to double when it grows at 1% (0.01) per period is given by,

$$T = 69.7 \simeq 70$$

At a growth rate of 2% the number of years is 69.3/2 ≃ 35, at a growth rate of 3% it is equal to 69.3/3 ≃ 23, at a growth rate of 5% it is equal to 69.3/5 ≃ 14, and at a growth rate of 10% it is equal to 69.3/10 ≃ 7.

Capital Accumulation and Economic Growth

In order to study the relation between savings, capital accumulation and economic growth we shall make two simplifying assumptions:

The first is that the size of the population, the participation rate, and the unemployment rate are all constant. This implies that employment, N, is also constant and equal to its "natural" rate.

The second assumption is that there is no technical progress, so the production function F does not change over time.

Then, in any period t, the relation between output per worker and capital per worker is given by,

$$(Y_t/N)=F(K_t/N\ ,\ 1)=f\ (K_t/N)$$

$f\ (K_t/N)$ is defined as being equal $F(K_t/N\ ,\ 1)$. In what follows we shall also denote output per worker as $y_t=Y_t/N$, and capital per worker as $k_t=K_t/N$. Hence, the aggregate production function can be written in the form,

$$y_t=f\ (k_t)$$

Hence, to the extent that there is capital accumulation per worker over time, output per worker also grows over time. This can be depicted diagrammatically in terms of a simple diagram. Suppose the economy at time 0 has capital k_0 per worker. Then its output per worker is equal to $y_0=f\ (k_0)$. If after T periods it has accumulated a higher capital stock, say $k_T > k_0$, then its output per worker will be equal to $y_T=f\ (k_T) > y_0=f\ (k_0)$.

Thus a higher capital stock per worker leads to higher output per worker over time.

Capital Accumulation and Economic Growth

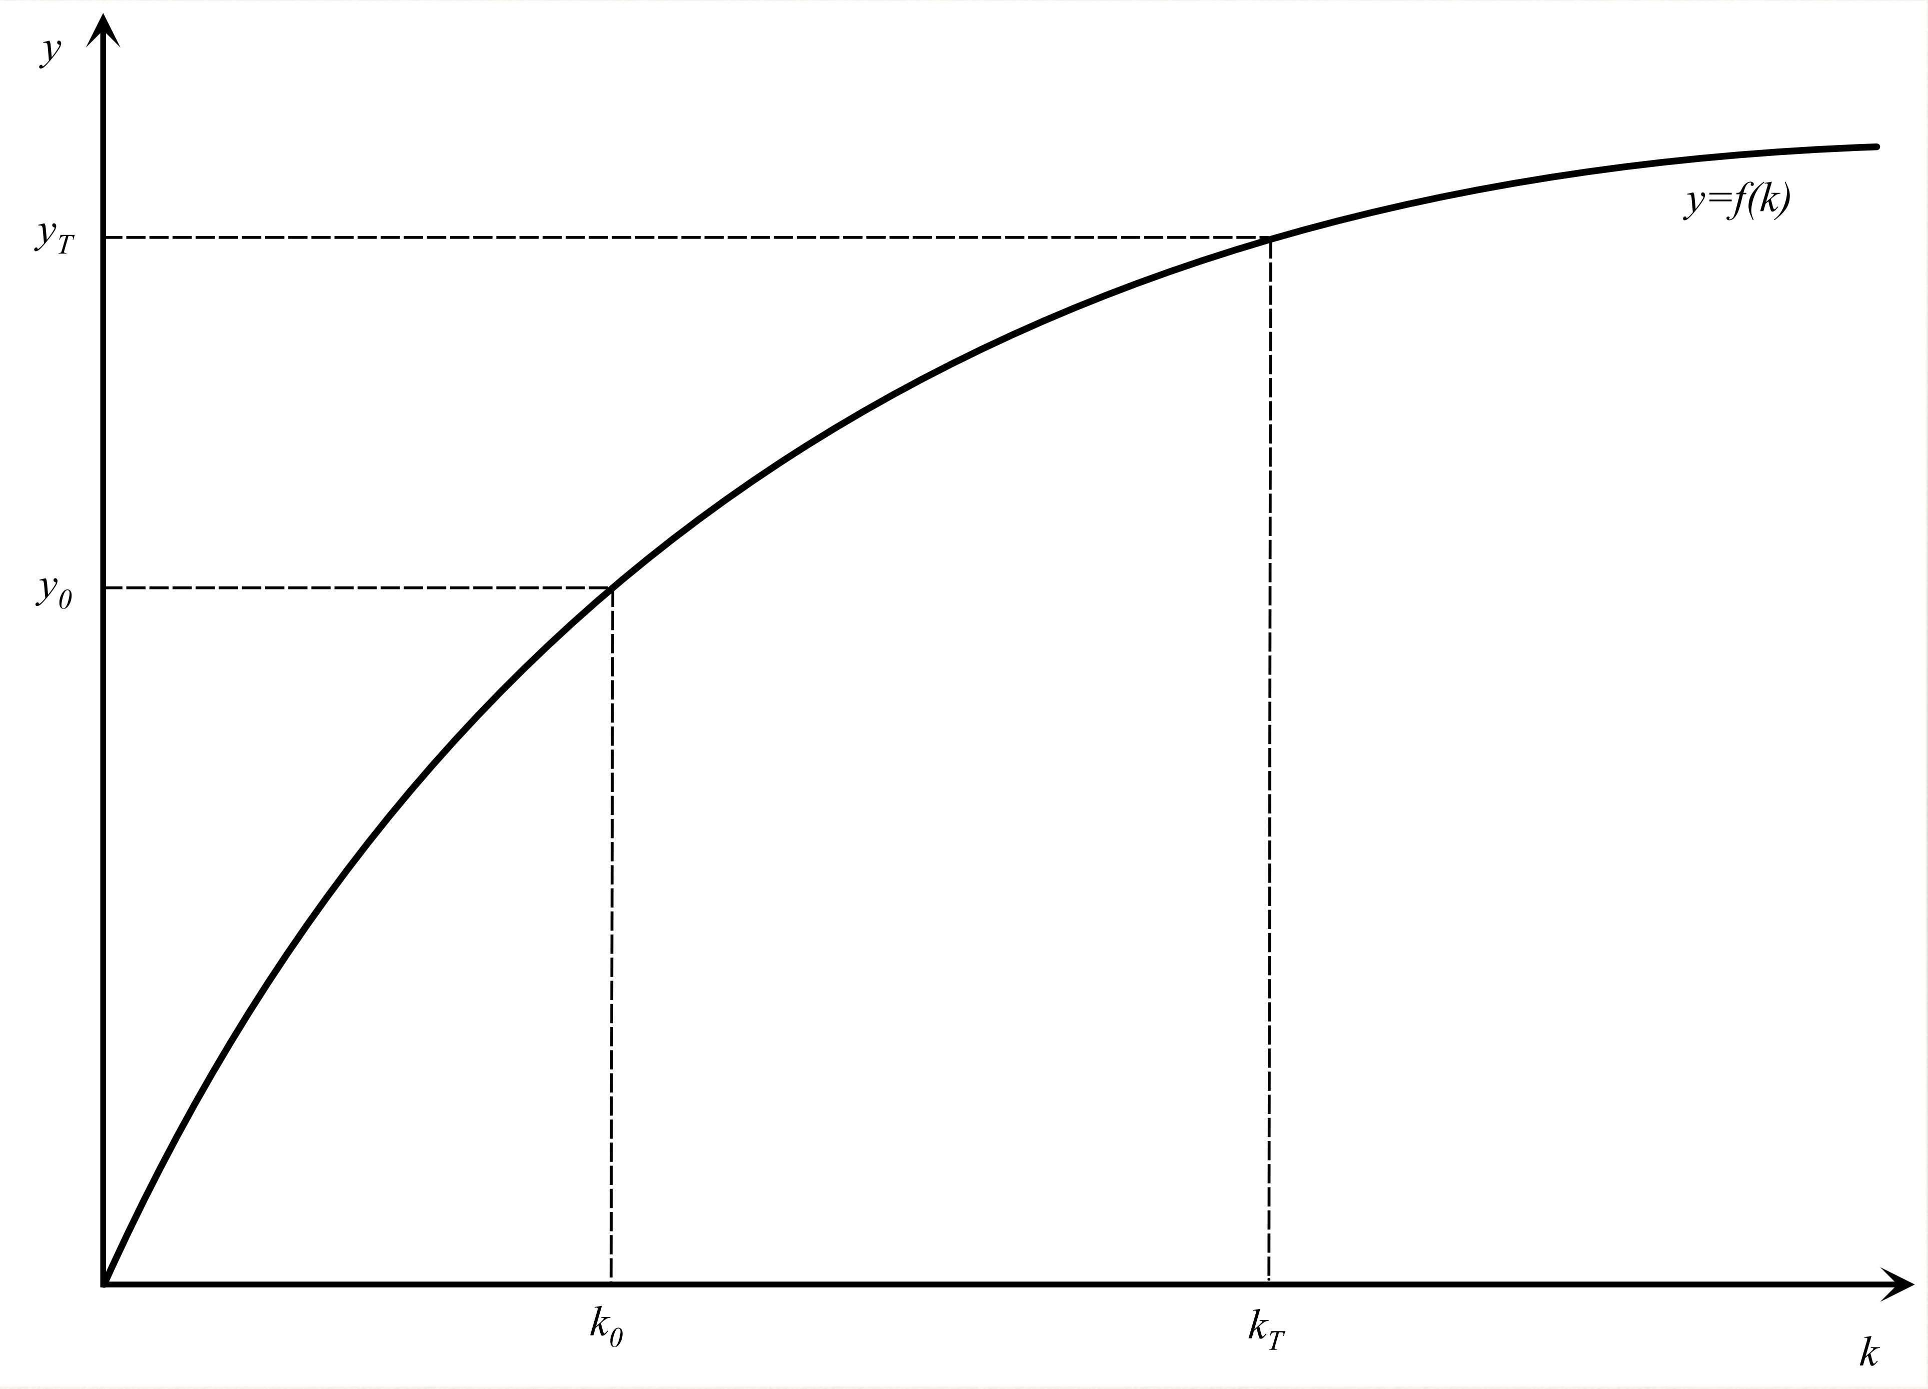

Savings, Investment and Capital Accumulation

The next question is how capital accumulates over time. It accumulates through investment. Investment is defined as *additions to the capital stock*, plus *replacement investment*, for capital that depreciates. If the depreciation rate δ is constant, the relation between capital accumulation and investment is given by,

$$I_t = K_{t+1} - K_t + \delta K_t$$

In per worker terms, investment is then given by,

$$k_{t+1} - k_t + \delta k_t$$

Recall that in a closed economy investment is equal to savings. Let us assume that savings per worker are proportional to income per worker. Thus, with a savings rate equal to s, savings per worker are given by,

$$sy_t = sf(k_t)$$

As a result, the accumulation of capital per worker is determined by,

$$k_{t+1} - k_t + \delta k_t = sf(k_t)$$

Hence, the accumulation of capital between period $t+1$ and period t is given by,

$$k_{t+1} - k_t = sf(k_t) - \delta k_t$$

To the extent that savings exceed the investment required for depreciation, capital accumulation is positive.

The Dynamics of Capital Accumulation

The accumulation of capital between period *t+1* and period *t* is equal to,

$$k_{t+1} - k_t = sf(k_t) - \delta k_t$$

To the extent that savings exceed the investment required for depreciation, capital accumulation is positive.

Solving the capital accumulation equation for the capital stock in period *t+1* allows us to study the relation between capital in each period relative to capital in the subsequent period. This is given by,

$$k_{t+1} = sf(k_t) + (1- \delta)k_t$$

Given the properties of the production function, it follows that capital per worker gradually converges to a constant level, irrespective of initial conditions. This constant level is called the long run equilibrium or *steady state* capital stock per worker. We shall denote it by k^*, which is the capital per worker which satisfies,

$$k^* = sf(k^*) + (1- \delta)k^*$$

The dynamics of capital accumulation and the relation between output growth and capital accumulation are depicted in the following two diagrams.

The Dynamics of Capital Accumulation

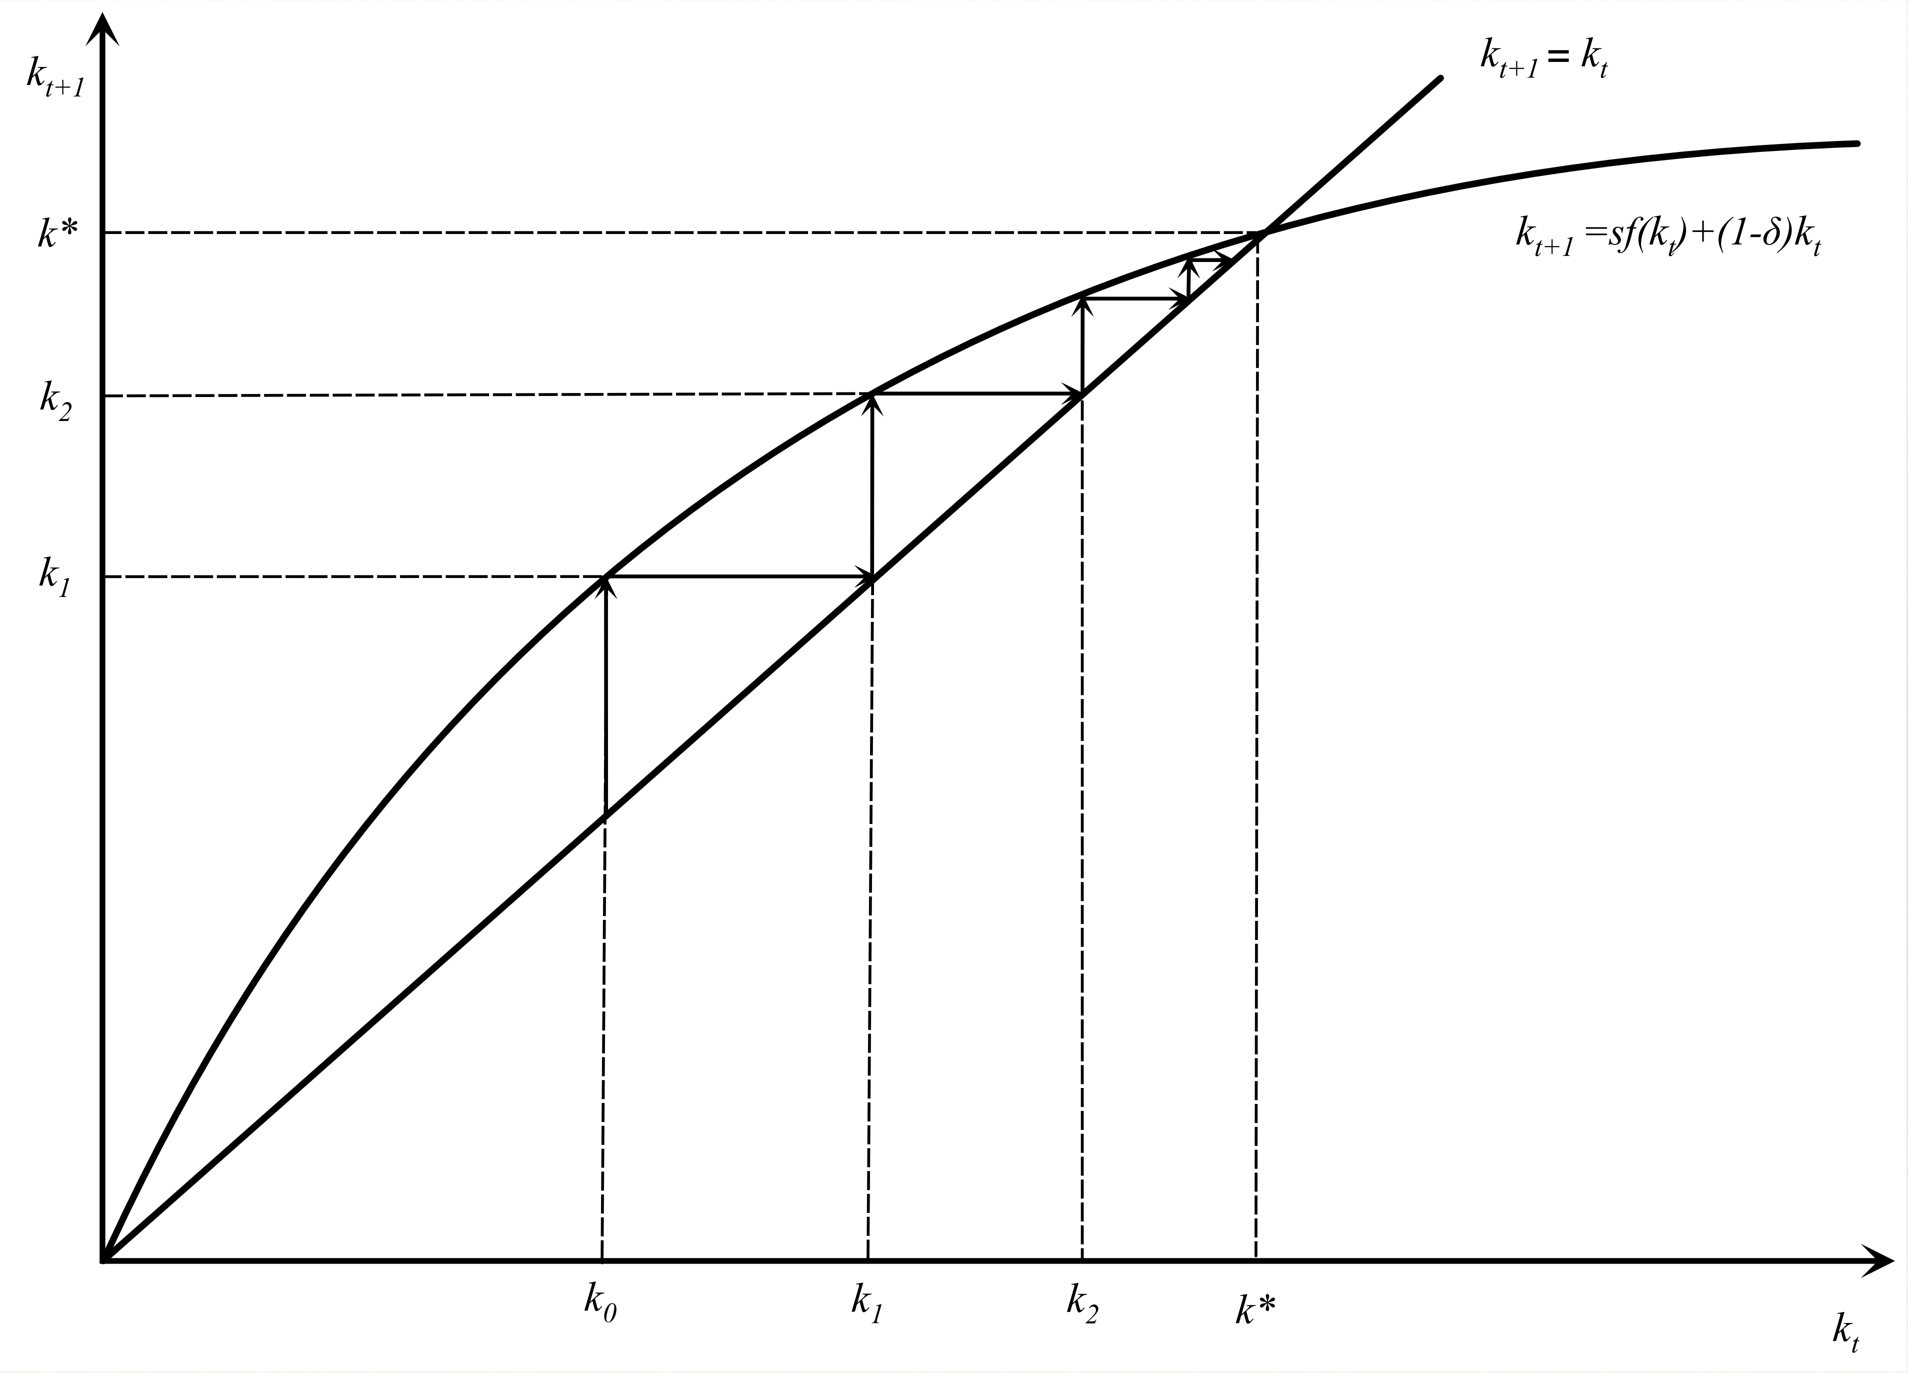

Savings, Capital Accumulation and Output Growth

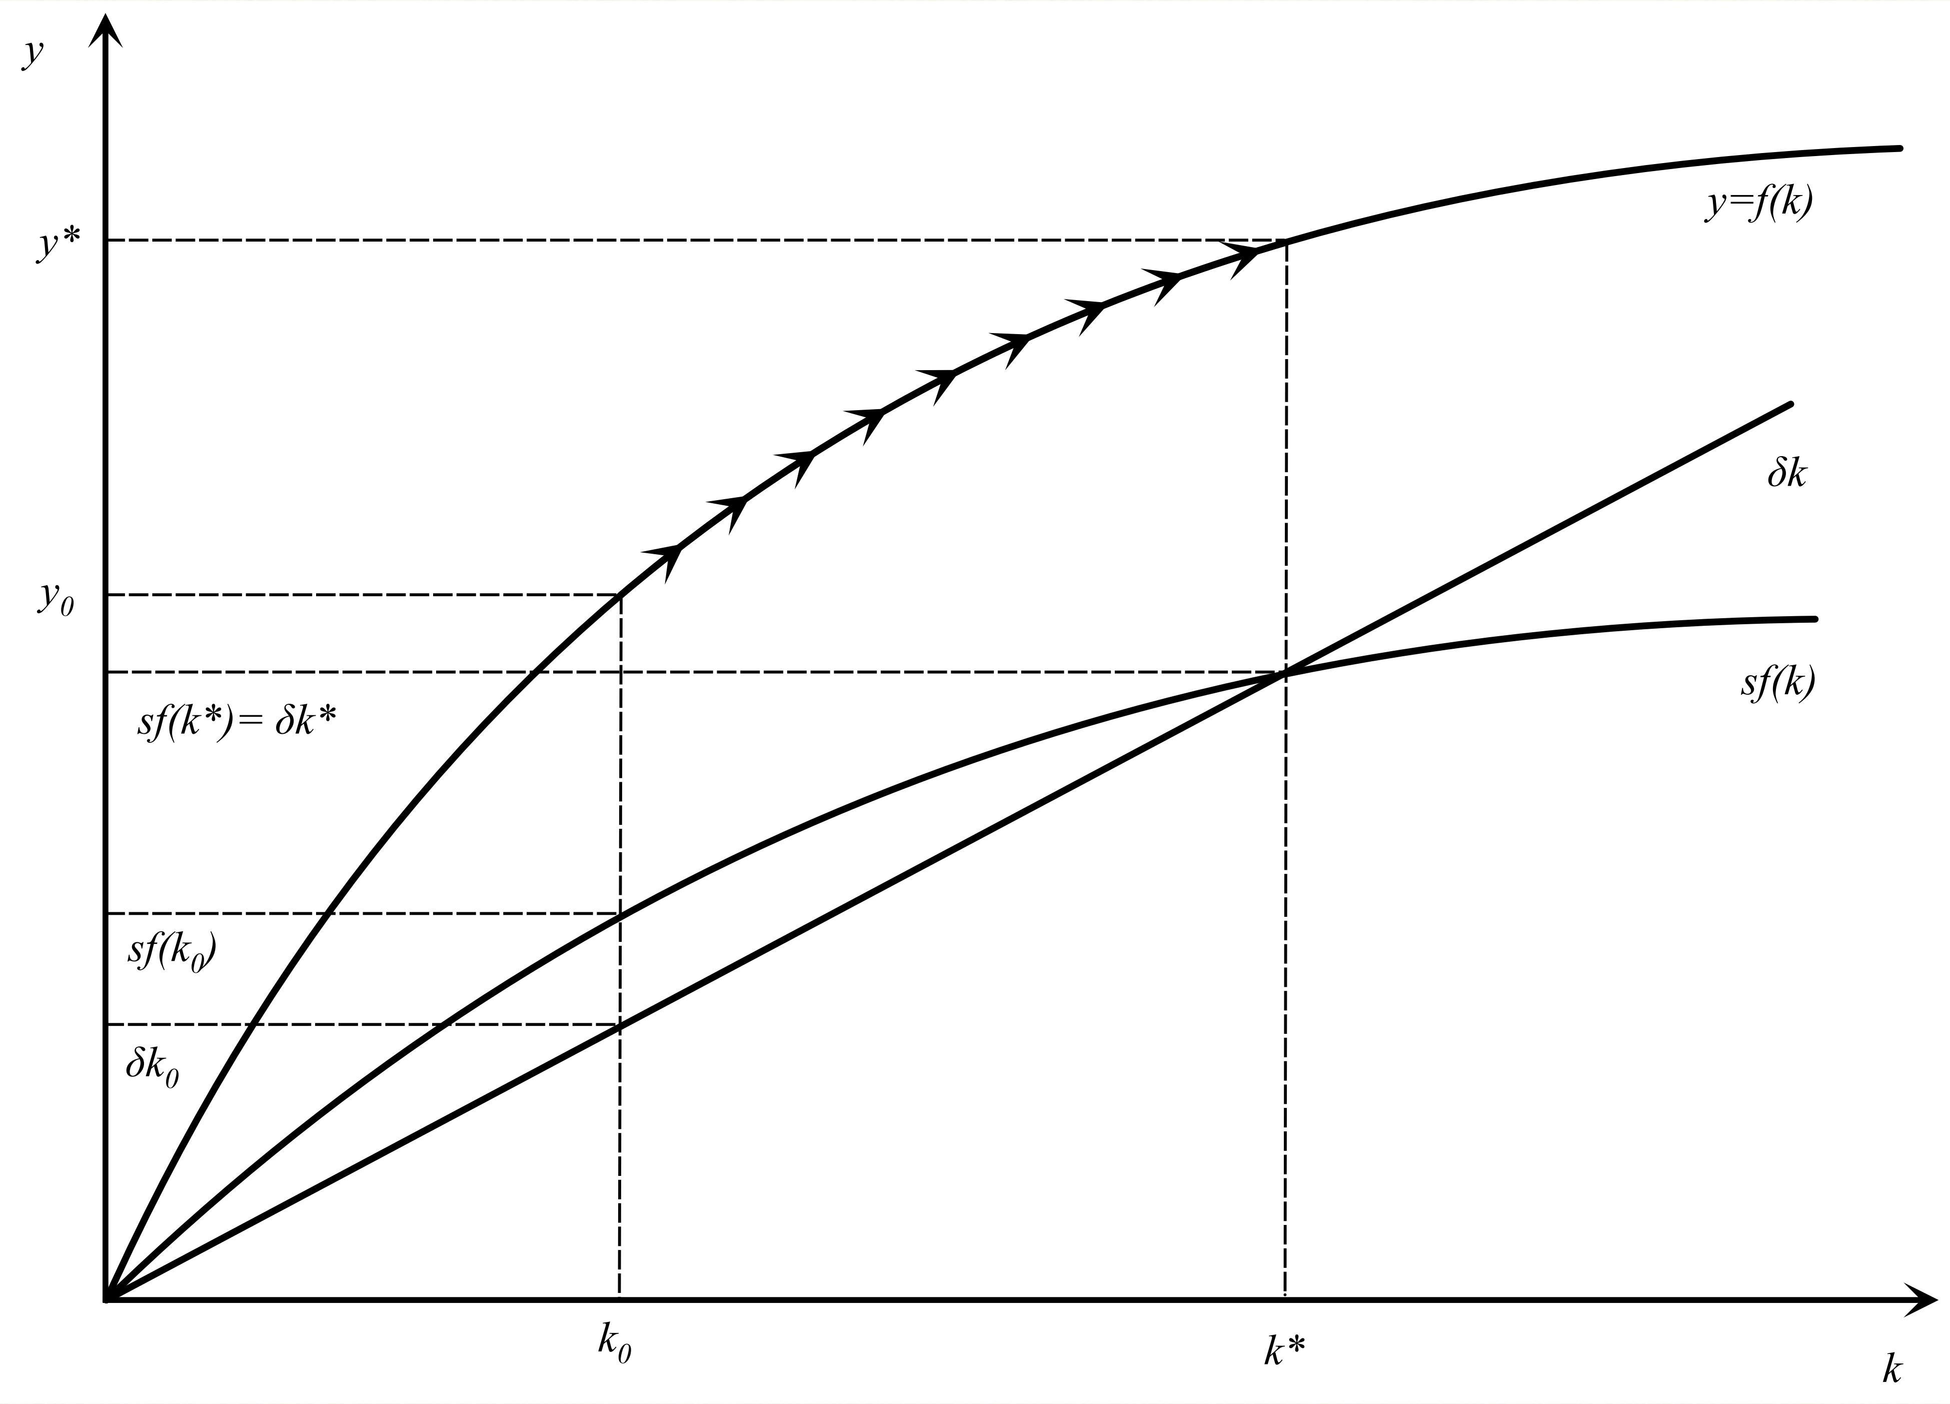

Steady State Output and the Savings Rate

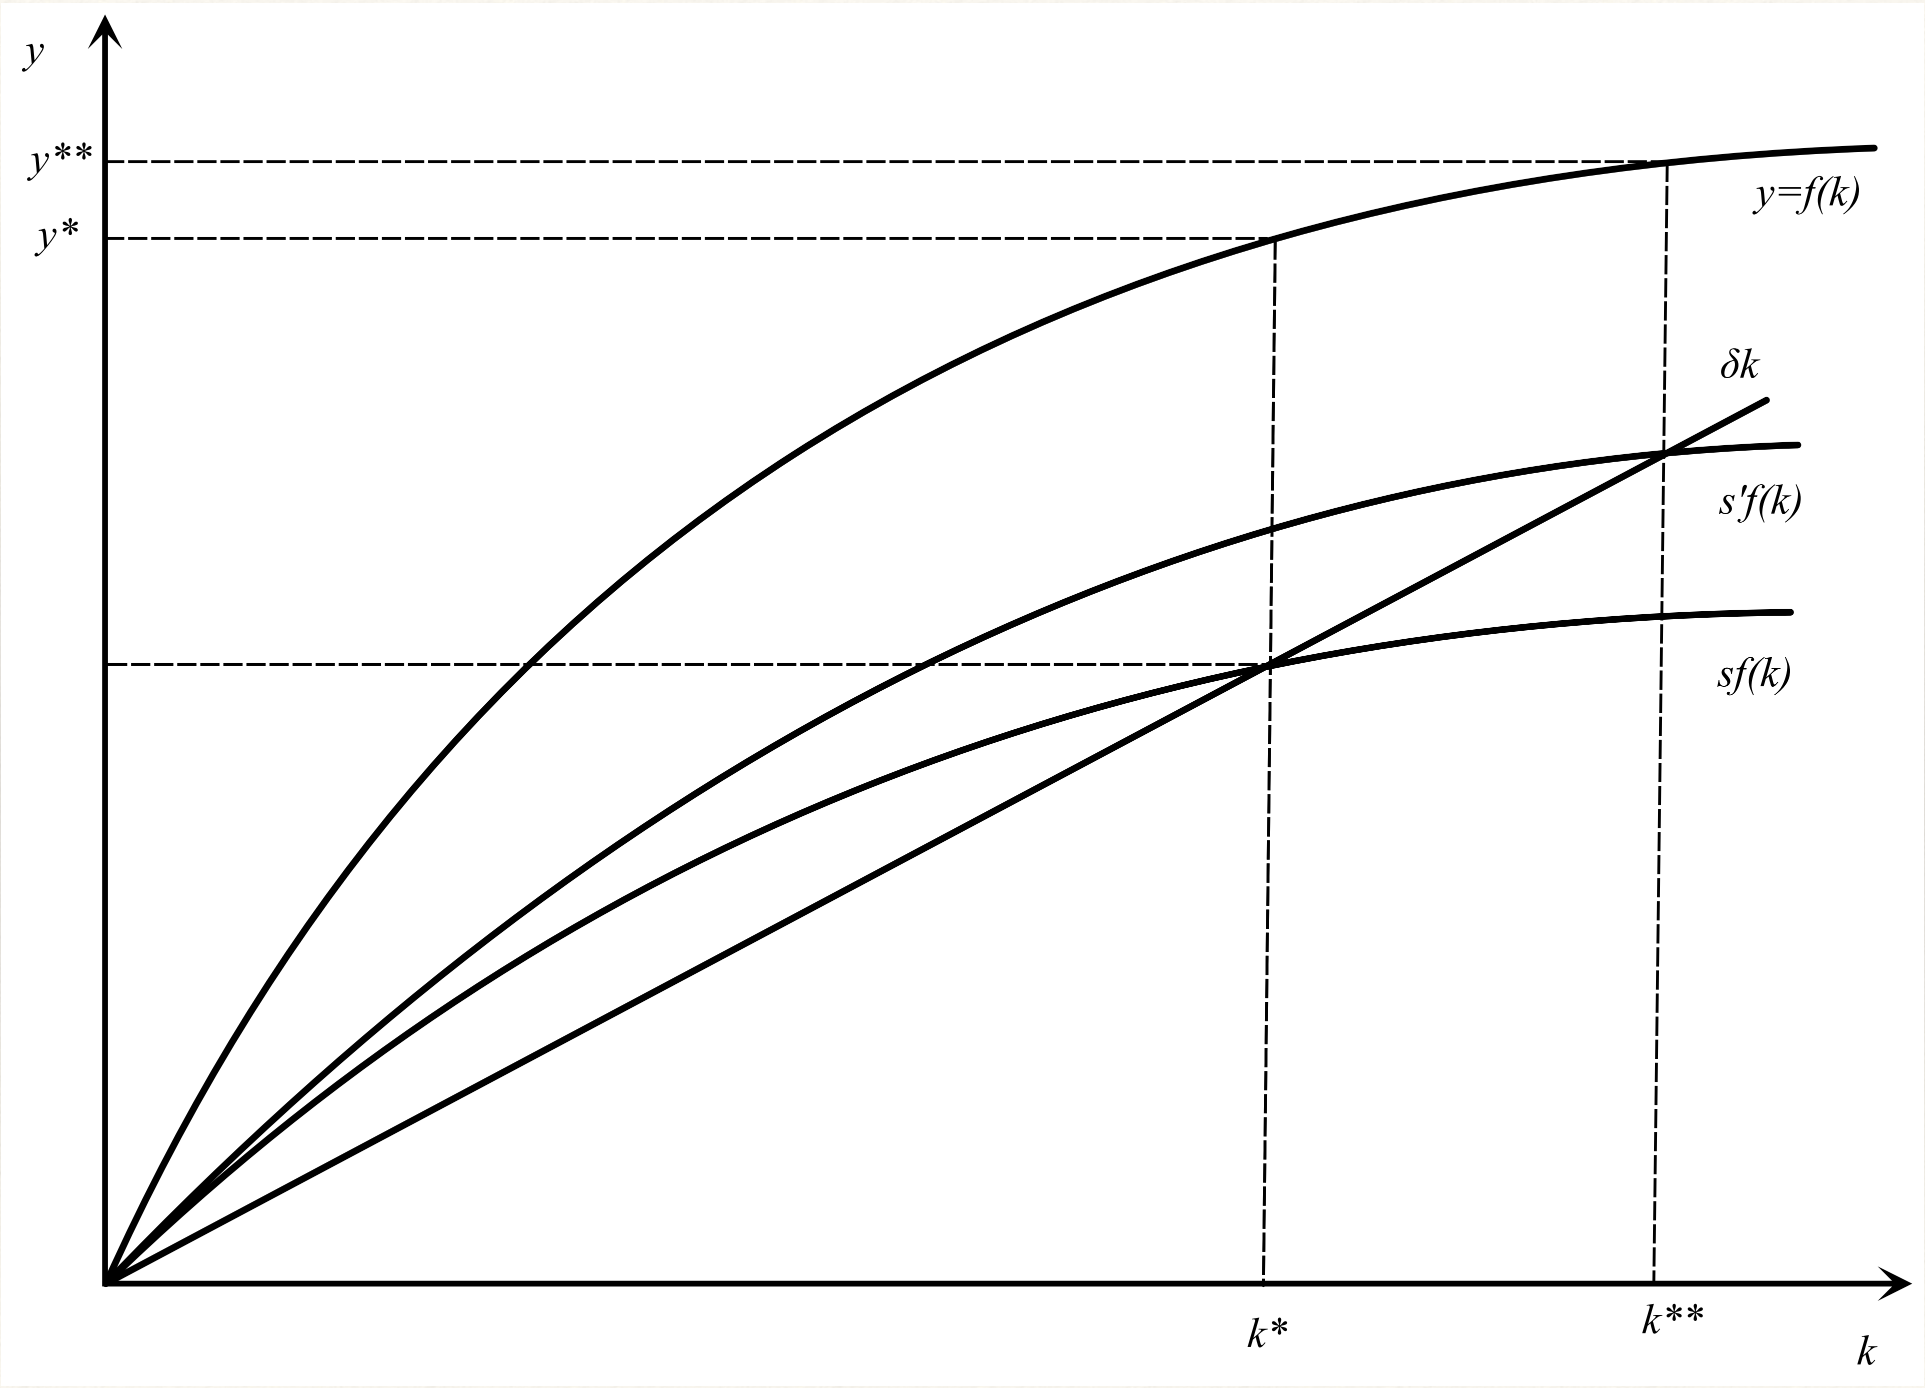

Transition to a Higher Steady State Output Path

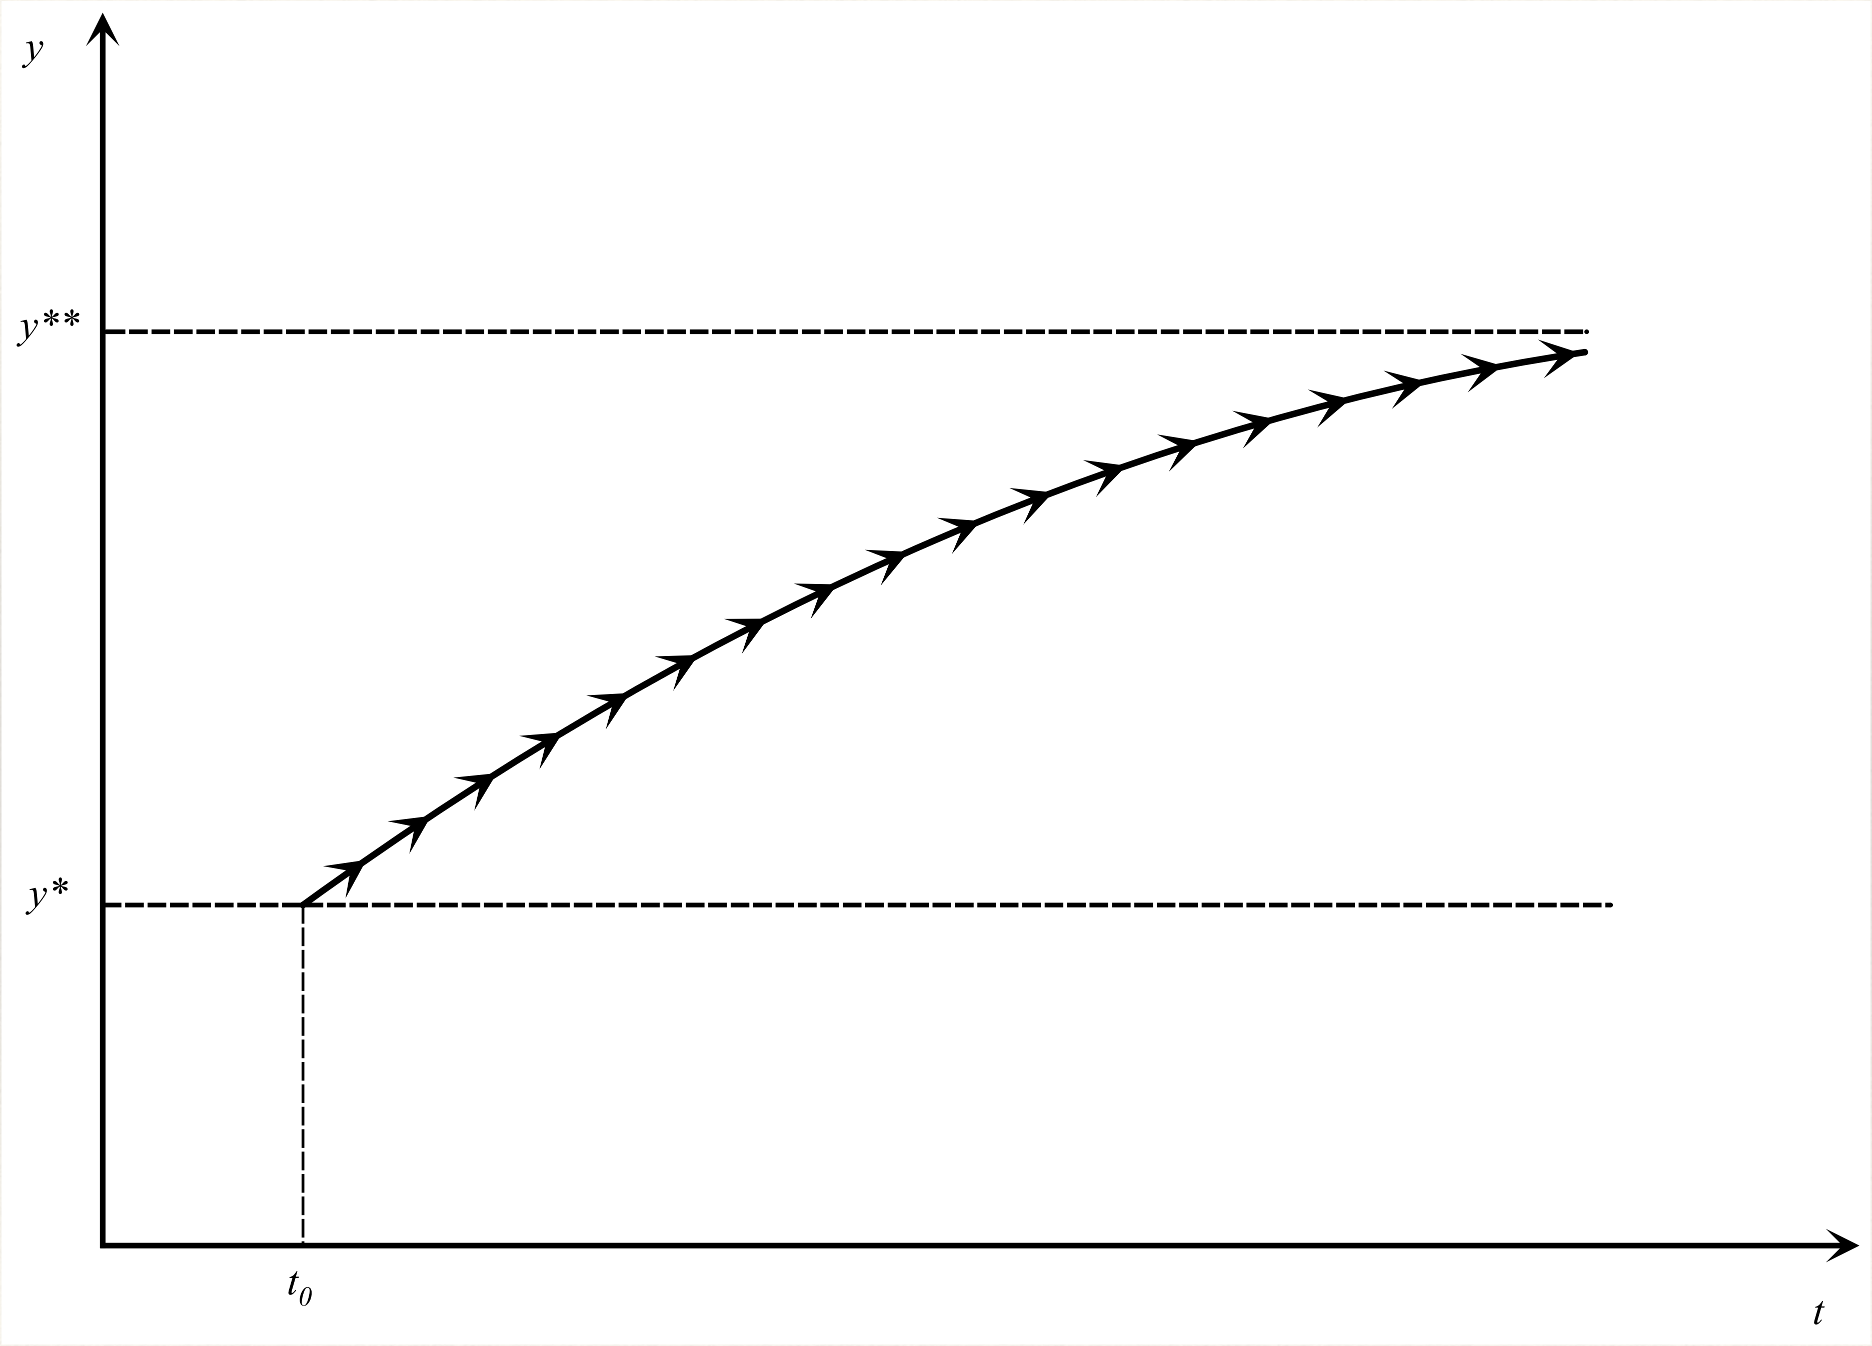

The Saving Rate and Consumption

Governments can affect the saving rate in various ways. First, they can vary public saving. Given private saving, positive public saving—a budget surplus, in other words—leads to higher overall saving. Conversely, negative public saving—a budget deficit—leads to lower overall saving. Second, governments can use taxes to affect private saving. For example, they can give tax breaks to people who save, making it more attractive to save and thus increasing private saving.

What saving rate should governments aim for? To think about the answer, we must shift our focus from the behavior of output to the behavior of consumption. The reason: What matters to people, especially in the long run, is not how much is produced, but how much they consume.

It is clear that an increase in saving must come initially at the expense of lower consumption. A change in the saving rate this year has no effect on capital this year, and consequently no effect on output and income this year. So an increase in saving comes initially with an equal decrease in consumption.

Does an increase in saving lead to an increase in consumption in the long run? Not necessarily. Consumption may decrease, not only initially, but also in the long run. To see why not, consider what happens for different values of the saving rate.

For a zero savings rate, capital and hence output and consumption are zero. For a savings rate of 100%, capital and hence output per worker is at its maximum, but consumption is also zero, because all income is saved, For savings rates in between consumption initially rises and after some point is starts falling.

The Savings Rate, Consumption and the Golden Rule

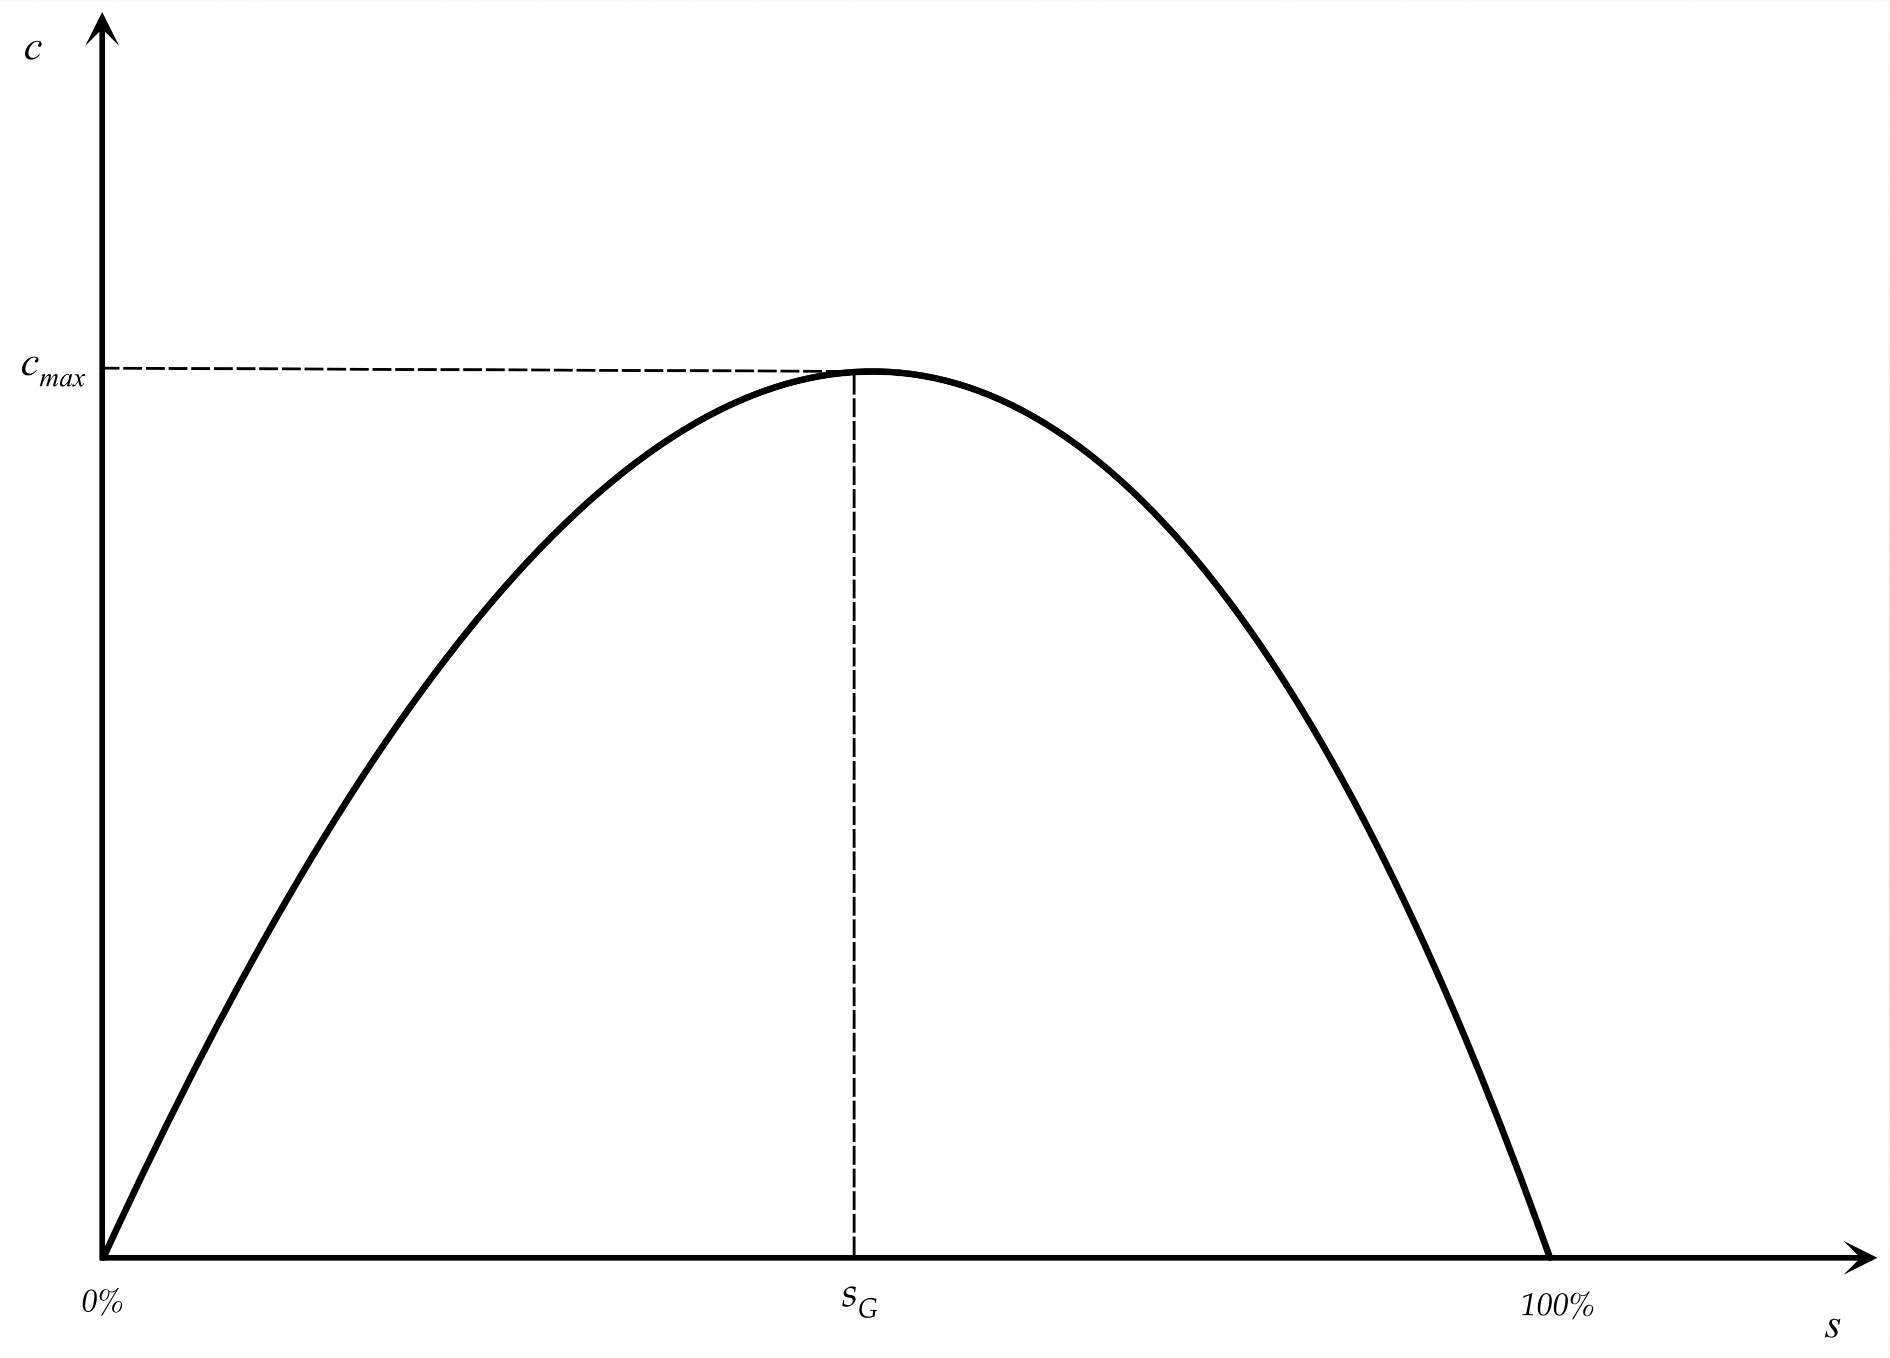

The Savings Rate, Consumption and the Golden Rule

A saving rate equal to zero implies a capital stock per worker equal to zero, a level of output per worker equal to zero, and, by implication, a level of consumption per worker equal to zero. For s between zero and s_G (G for golden rule), a higher saving rate leads to higher capital per worker, higher output per worker, and higher consumption per worker. For s larger than s_G, increases in the saving rate still lead to higher values of capital per worker and output per worker; but they now lead to lower values of consumption per worker: This is because the increase in output is more than offset by the increase in depreciation due to the larger capital stock. For $s = 1$, consumption per worker is equal to zero. Capital per worker and output per worker are high, but all of the output is used just to pay for depreciation, leaving nothing for consumption.

If an economy already has so much capital that it is operating beyond the golden rule, then increasing saving further will decrease consumption not only now, but also later (in the steady state). Is this a relevant worry? Do some countries actually have too much capital? The empirical evidence indicates that most OECD countries are actually far below their golden-rule level of capital. If they were to increase the saving rate, it would lead to higher consumption in the future—not lower consumption.

This means that, in practice, governments face a trade-off: An increase in the saving rate leads to lower consumption for some time, but higher consumption later. So what should governments do? How close to the golden rule should they try to get? That depends on how much weight they put on the welfare of current generations—who are more likely to lose from policies aimed at increasing the saving rate—versus the welfare of future generations—who are more likely to gain. Enter politics: Future generations do not vote. This means that governments are unlikely to ask current generations to make large sacrifices, which, in turn, means that capital is likely to stay far below its golden-rule level.

Properties of the Solow Model

The model we have just presented was first put forward by Robert Solow in 1956. It is thus known as the Solow model.

According to the model, in the long run, the evolution of output per worker is determined by two relations. First, the level of output per worker depends positively on the amount of capital per worker, through a neoclassical production function. Second, capital accumulation depends on the level of output, which determines saving and investment. It is assumed that investment is a constant share of output, determined by the savings rate.

These interactions between capital and output imply that, starting from any level of capital, in the absence of technical progress, an economy converges in the long run to a steady-state (constant) level of capital. Associated with this level of capital is a *steady-state level of output*.

1. The steady-state level of capital per worker, and thus the steady-state level of output per worker, depends positively on the savings rate. A higher saving rate leads to a higher steady-state level of capital and output.

2. During the transition to the new steady state, a higher saving rate leads to positive output growth. But in the absence of technical progress, in the long run, the growth rate of output is equal to zero and so does not depend on the saving rate.

3. In the absence of technical progress, the Solow model cannot explain positive output growth in the steady state, but only the level of per capita income.

Robert M. Solow (1924 -)

Robert Merton Solow, (born August 23, 1924), is a leading American economist, particularly known for his work on the theory of economic growth that culminated in the growth model named after him.

He is currently Emeritus Institute Professor of Economics at the Massachusetts Institute of Technology, where he has been a professor since 1949. He was awarded the John Bates Clark Medal of the American Economic Association in 1961, the Nobel Memorial Prize in Economic Sciences in 1987, and the Presidential Medal of Freedom in 2014.

Physical versus Human Capital

We have concentrated so far on *physical capital*—machines, plants, office buildings, and so on. But economies have another type of capital: the set of skills of the workers in the economy, or what economists call *human capital*. An economy with many highly skilled workers is likely to be much more productive than an economy in which most workers cannot read or write.

The increase in human capital has been as large as the increase in physical capital over the last two centuries. At the beginning of the Industrial Revolution, only 30% of the population of the countries that constitute the OECD today knew how to read. Today, the literacy rate in OECD countries is above 95%. Schooling was not compulsory prior to the Industrial Revolution. Today it is compulsory, usually until the age of 16. Still, there are large differences across countries. Today, in OECD countries, nearly 100% of children get a primary education, 90% get a secondary education, and 38% get a higher education. The corresponding numbers in poor countries, countries with GDP per person below $400, are 95%, 32%, and 4%, respectively.

How should we think about the effect of human capital on output? How does the introduction of human capital change our earlier conclusions?

Extending the Production Function to Allow for Human Capital

The most natural way of extending our analysis to allow for human capital is to modify the production function relation to read,

$$y = f(k, h)$$

where is $y=Y/N$ output per worker, $k=K/N$ is physical capital per worker and $h=H/N$ is human capital per worker.

As before, an increase in capital per worker k leads to an increase in output per worker. And an increase in the average level of skill h also leads to more output per worker. More skilled workers can do more complex tasks; they can deal more easily with unexpected complications. All of this leads to higher output per worker.

We assumed earlier that increases in physical capital per worker increased output per worker, but that the effect became smaller as the level of capital per worker increased. We can make the same assumption for human capital per worker: Think of increases in h as coming from increases in the number of years of education and training.

Human Capital, Physical Capital and Output

How does the introduction of human capital change the analysis of the previous sections? Our conclusions about physical capital accumulation remain valid: An increase in the saving rate increases steady-state physical capital per worker and therefore increases output per worker.

But our conclusions now extend to human capital accumulation as well. An increase in how much society "saves" in the form of human capital—through education and on-the-job training—increases steady-state human capital per worker, which leads to an increase in output per worker. Our extended model gives us a richer picture of how output per worker is determined. In the long run, it tells us that output per worker depends on both how much society saves and how much it spends on education.

What are the relative contributions of human capital and physical capital in the determination of output per worker? A place to start is to compare how much is spent on formal education to how much is invested in physical capital.

In the United States, spending on formal education is about 6.5% of GDP. This number includes both government expenditures on education and private expenditures by people on education. It is between one-third and one-half of the gross investment rate for physical capital (which is around 16%). But this comparison is only a first pass, as education is also partly consumption, the costs of acquiring an education include foregone wages, a lot of human capital is acquired through on the job training, and the depreciation rate of human capital may differ from physical capital. Hence it is more difficult to come up with reliable numbers for the accumulation of human capital from spending on education alone.

Human Capital, Physical Capital and Output

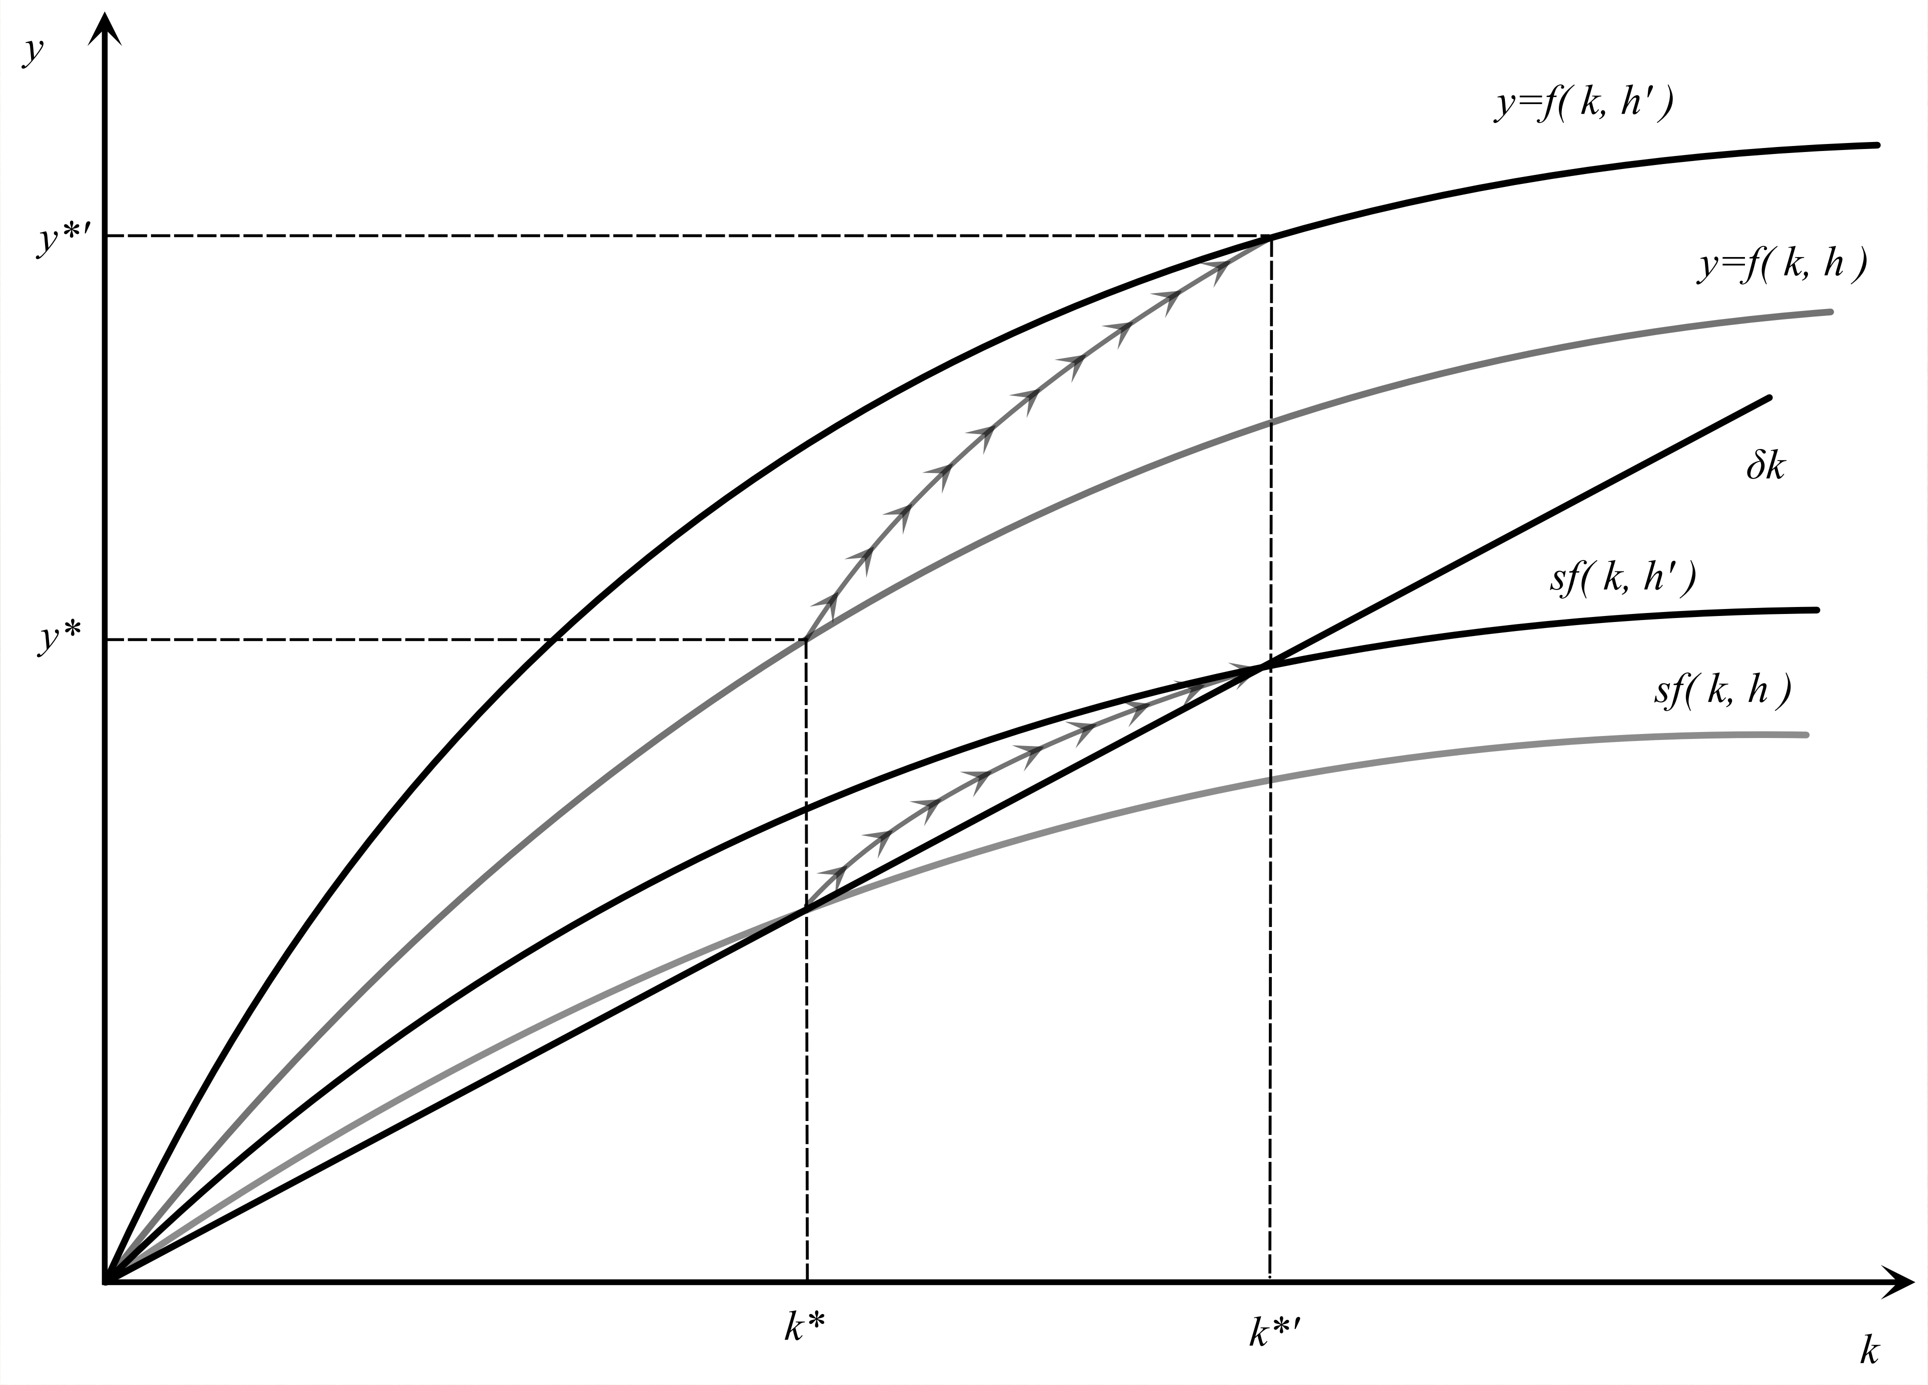

Human Capital Accumulation and Long Term Growth

Even in a Solow model with human capital accumulation, a country that saves more or spends more on education will achieve a higher level of output per worker in steady state. It does not follow that by saving or spending more on education and training a country can sustain permanently higher growth of output per worker.

For this, human and physical capital must accumulate jointly, in such a way that we never reach a steady state. Given human capital, increases in physical capital will run into decreasing returns. And given physical capital, increases in human capital will also run into decreasing returns. But what if both physical and human capital increase in tandem? Can't an economy grow forever just by steadily having more capital and more skills per worker?

Models that generate steady growth even without technological progress are called *models of endogenous growth* to reflect the fact that in those models - in contrast to the Solow model - the growth rate depends, even in the long run, on variables such as the saving rate and the rate of spending on education.

The empirical evidence suggests that physical and human capital accumulation do not suffice to explain long term growth. Thus, a wide consensus among economists is a follows:

Output per worker depends on the level of both physical capital per worker and human capital per worker. Both forms of capital can be accumulated, one through physical investment, the other through education and training. Increasing either the saving rate and/or the fraction of output spent on education and training can lead to much higher levels of output per worker in the long run. However, such measures do not lead to a permanently higher growth rate. One needs to assume technical progress in order to explain long term growth.

Conclusions About Physical and Human Capital Accumulation and Economic Growth

In the long run, the evolution of output per worker is determined by two relations. First, on the amount of capital per worker. Second, on savings rate and the level of output per worker, which determines capital accumulation.

These interactions between capital and output imply that, starting from any level of capital per worker, in the absence of technical progress, an economy converges in the long run to a steady-state (constant) level of capital per worker. Associated with this level of capital per worker is a steady-state level of output per worker.

The steady-state level of capital per worker, and thus the steady-state level of output per worker, depends positively on the saving rate. A higher saving rate leads to a higher steady-state level of output per worker. During the transition to the new steady state, a higher saving rate leads to positive output growth. But, in the long run, in the absence of technical progress, the growth rate of output per worker is equal to zero and so does not depend on the saving rate.

These conclusions are not modified by allowing for human capital accumulation through education and training, unless the accumulation of human capital results in non decreasing returns from physical capital accumulation and there is endogenous growth.

The Nature of Technical Progress

Technical progress has many dimensions:

It can lead to larger quantities of output for given quantities of capital and labor: Think of a new type of machine that runs at a higher speed, and leads to increased production.

It can lead to better products: Think of the steady improvement in automobile safety and comfort over time, or better smartphones.

It can lead to new products: Think of the introduction of the CD or MP3 player, the smartphone, wireless communication technology in all its variants, flat screen monitors, and high-definition television.

It can lead to a larger variety of products: Think of the steady increase in the number of breakfast cereals available at your local supermarket.

These dimensions are more similar than they appear. If we think of consumers as caring not about the goods themselves but about the services these goods provide, then they all have something in common: In each case, consumers receive more services. A better car provides more safety, a new product such as the smart phone or a new service such as wireless communication technology provides more communication services, and so on. If we think of output as the set of underlying services provided by the goods produced in the economy, we can think of technical progress as leading to increases in output for given amounts of capital and labor. We can then think of the state of technology as a variable that tells us how much output can be produced from given amounts of capital and labor at any time.

The State of Technology and the Production Function

If we denote the state of technology by A, we can rewrite the production function as,

$$Y = F(K, AN)$$

This equation states that production depends on capital and on labor multiplied by the state of technology. The state of technology determines the efficiency of labor.

This way of introducing the state of technology implies that we can think of technical progress in two equivalent ways:

1. Technical progress reduces the number of workers needed to produce a given amount of output. Doubling A produces the same quantity of output with only half the original number of workers, N.
2. Technical progress increases the output that can be produced with a given number of workers. We can think of AN as the amount of *effective labor* in the economy. If the state of technology A doubles, it is as if the economy had twice as many workers. In other words, we can think of output being produced by two factors: capital (K), and effective labor (AN).

If we again assume *constant returns to scale*, we can rewrite the production function as,

$$y = f(k)$$

Where $y=Y/AN$ is output per effective unit of labor and $k=K/AN$ is capital per effective unit of labor.

The Neoclassical Production Function
Output per Effective Unit of Labor

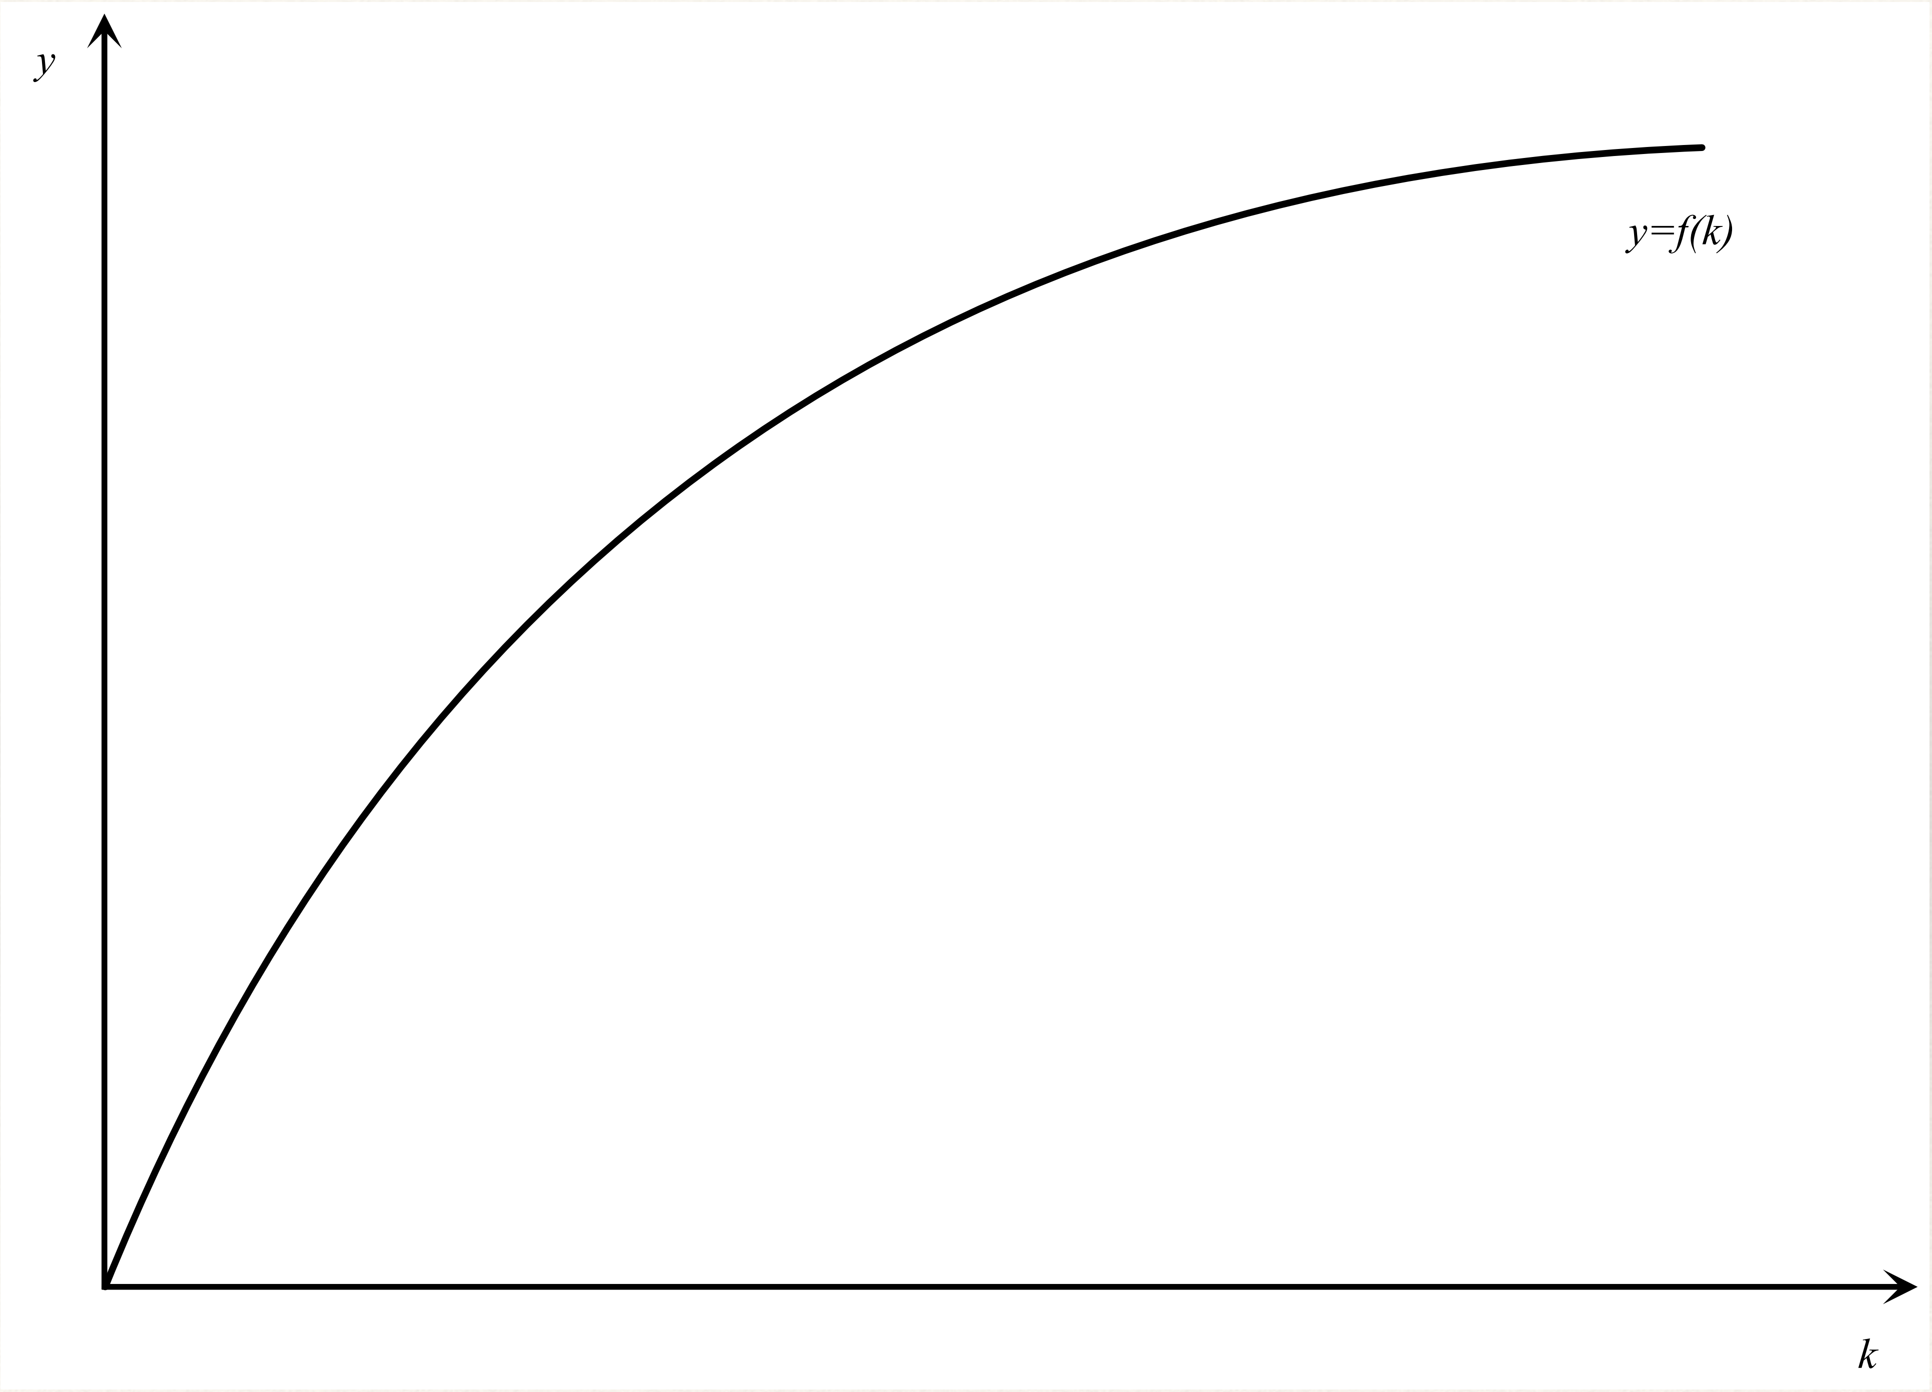

Population Growth and Exogenous Technical Progress

Assume that the number of workers (population) grows over time at an exogenous rate $n > 0$, the rate of *population growth*. Hence,

$$N_t=(1+n)N_{t-1}$$

Also assume that the efficiency of labor A grows over time at an exogenous rate $g > 0$, the rate of *technical progress*. Hence,

$$A_t=(1+g)A_{t-1}$$

Together these assumptions imply that the growth of effective labor is given by $g+n$.

Thus, if $n = 1\%$ per annum, and $g = 2\%$ per annum, the rate of growth of effective labor is equal to 3% per annum.

Savings, Investment and Capital Accumulation in the Presence of Population Growth and Technical Progress

The next question is how capital accumulates over time. It accumulates through investment. Investment is defined as *additions to the capital stock*, plus *replacement investment*, for capital that depreciates. If the depreciation rate δ is constant, the relation between capital accumulation and investment is given by,

$$I_t = K_{t+1} - K_t + \delta K_t$$

Dividing through by effective labor, investment per effective unit of labor is then given by,

$$(1+g)(1+n)k_{t+1} - k_t + \delta k_t$$

Recall that in a closed economy investment is equal to savings. As in the model without population growth and technical progress, we assume that savings per effective unit of labor are proportional to income per worker. Thus, with a savings rate equal to s, savings per worker are given by,

$$sy_t = sf(k_t)$$

As a result, the accumulation of capital per effective unit of labor is determined by the savings investment relation,

$$(1+g)(1+n)k_{t+1} - k_t + \delta k_t = sf(k_t)$$

To the extent that savings exceed the investment required for technical progress, population growth and depreciation, capital accumulation per effective unit of labor is positive.

The Dynamics of Capital Accumulation

For relatively small g and n (less than 5-10%), the accumulation of capital between period $t+1$ and period t can be shown to be approximately equal to,

$$k_{t+1} - k_t = sf(k_t) - (g+n+\delta)k_t$$

To the extent that savings exceed the investment required for depreciation, technical progress and population growth, capital accumulation per effective unit of labor is positive.

Solving the capital accumulation equation for the capital stock in period $t+1$ allows us to study the relation between capital in each period relative to capital in the subsequent period. This is given by,

$$k_{t+1} = sf(k_t) +(1-g-n-\delta)k_t$$

Given the properties of the production function, it follows that capital per effective unit of labor gradually converges to a constant level, irrespective of initial conditions. This constant level is called the long run equilibrium or *steady state* capital stock per effective unit of labor. We shall denote it by k^*, which is the capital per effective unit of labor which satisfies,

$$k^* = sf(k^*) +(1- \delta)k^*$$

The dynamics of capital accumulation and the relation between output growth and capital accumulation are depicted in the following two diagrams, similar to the diagrams of an economy without population growth and technical progress.

The Dynamics of Capital Accumulation in the Presence of Technical Progress and Population Growth

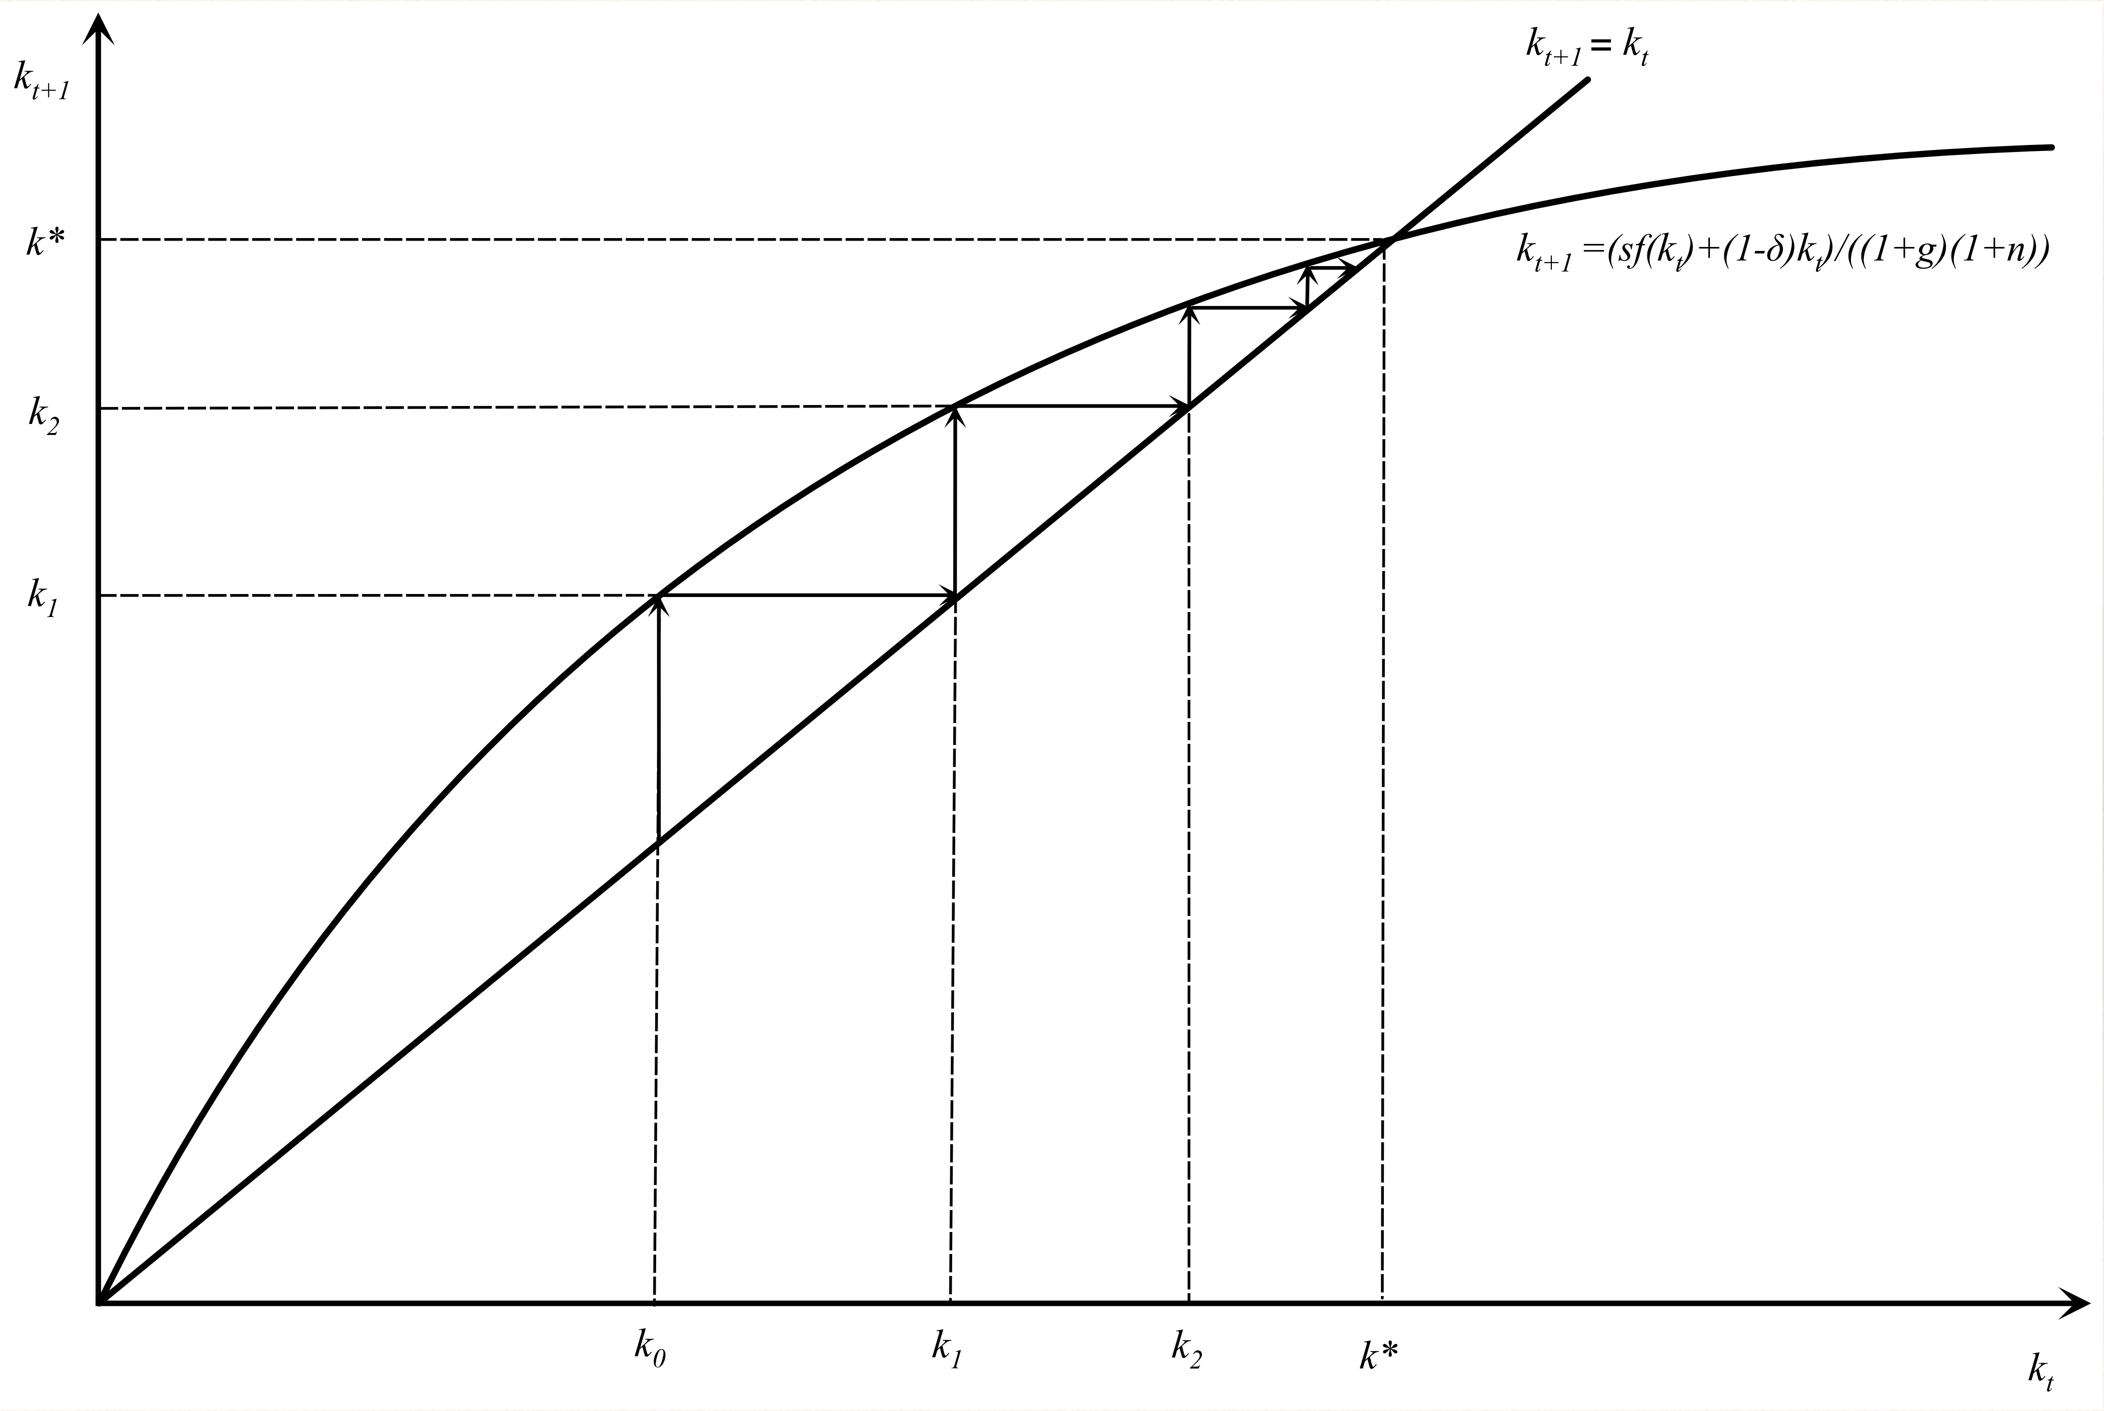

Savings, Capital Accumulation and Output Growth

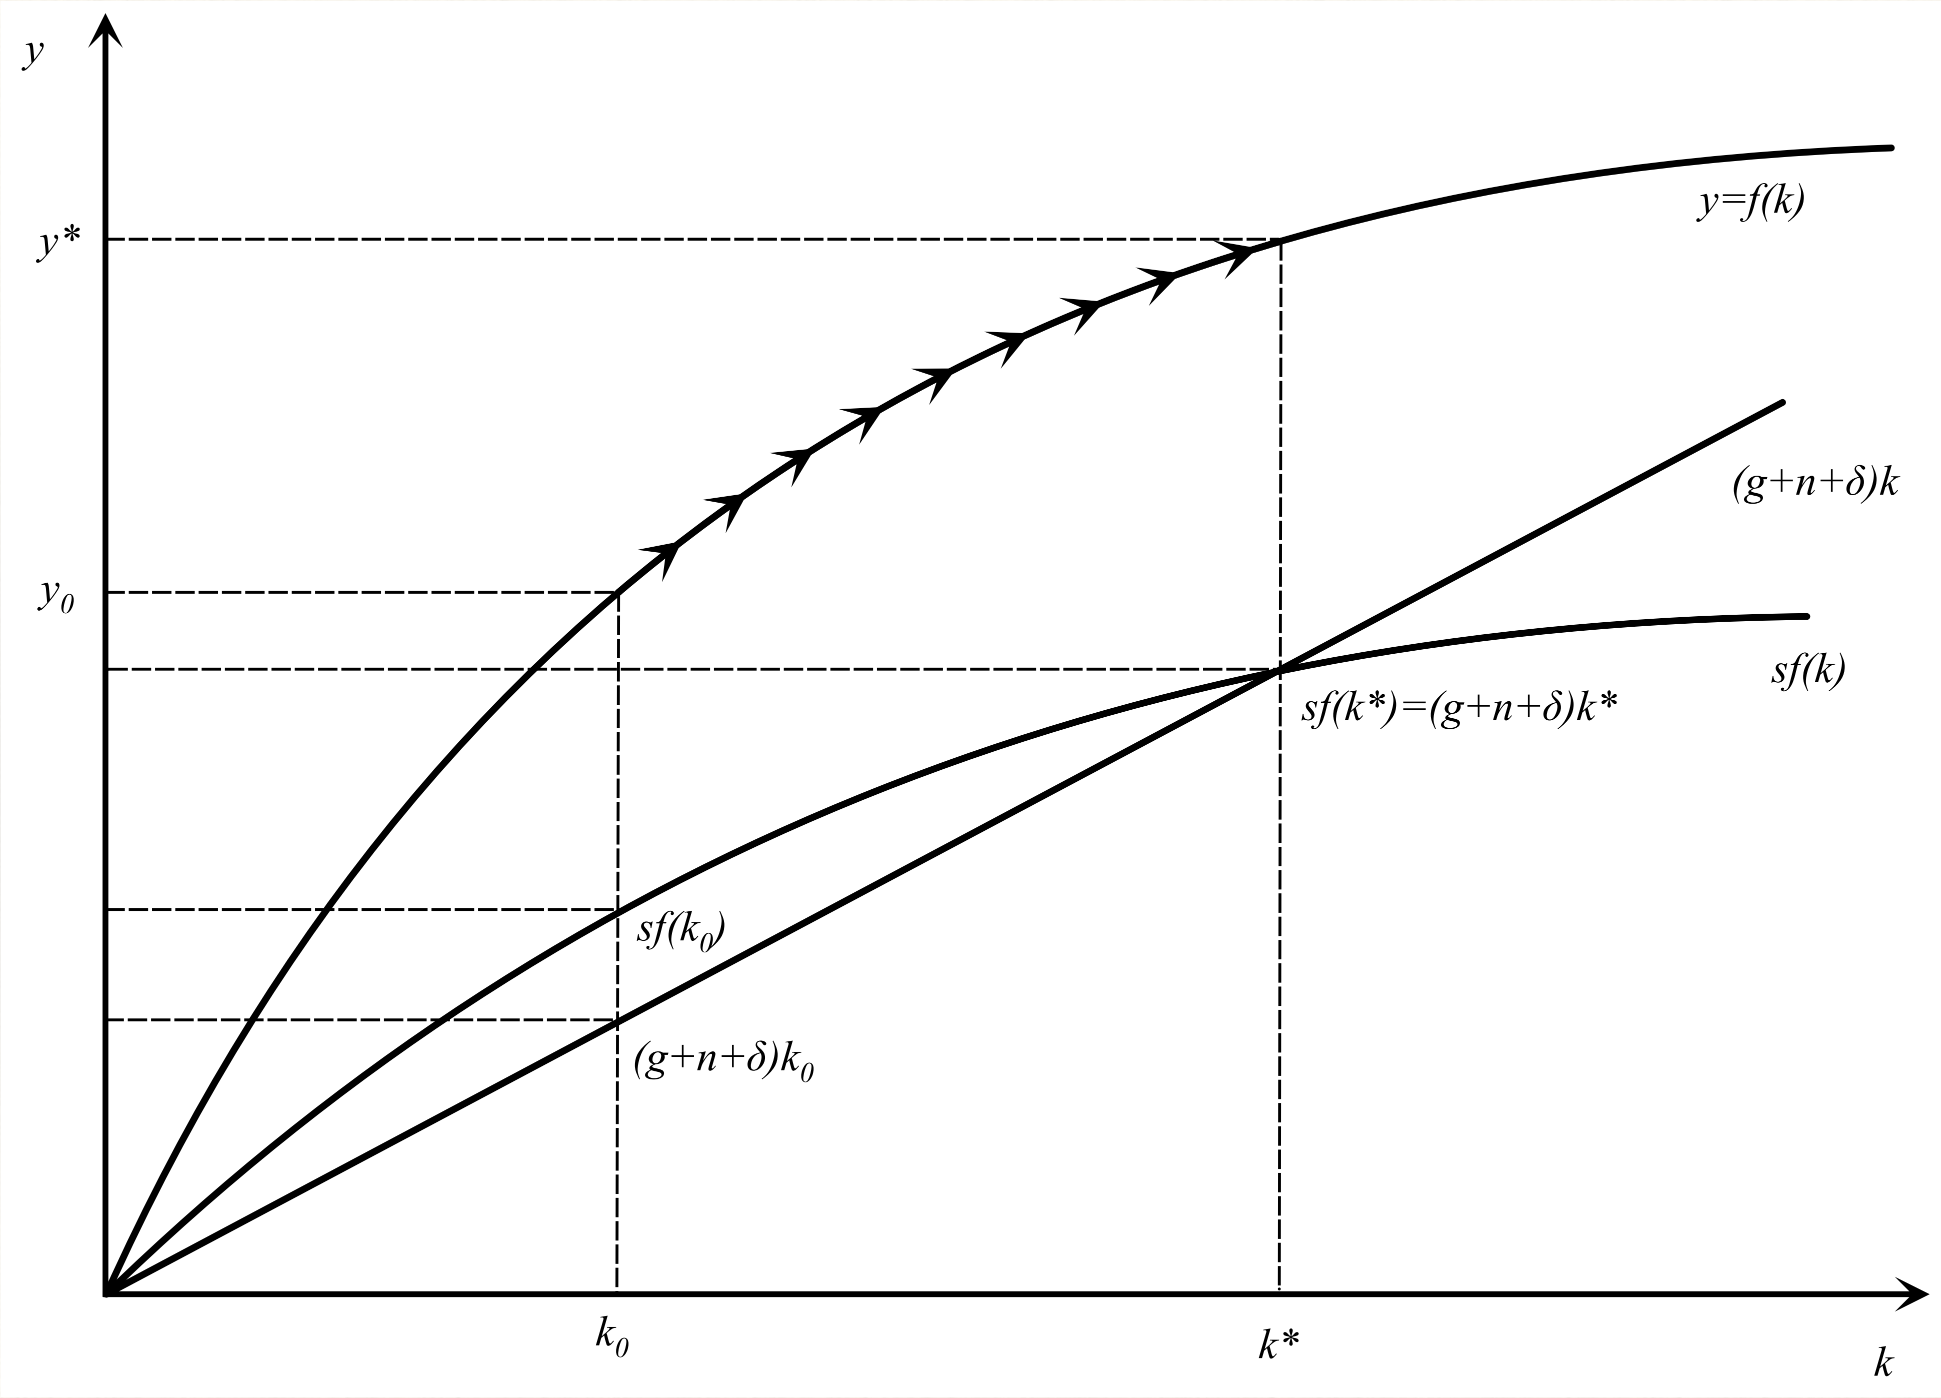

The Properties of the Steady State in the Presence of Technical Progress and Population Growth

In the long run, capital per effective unit of labor reaches a constant level, and so does output per effective unit of labor. Put another way, the steady state of this economy is such that capital and output per effective unit of labor are constant and equal $k^*=(K_t/(A_tN_t)^*$ and $y^*=(Y_t/(A_tN_t)^*$, respectively.

How about aggregate capital and output. Both grow at the sum of the rate of technical progress g, and population growth n, since, in steady state, capital and output are given by,

$$K_t=k^*A_tN_t$$

$$Y_t=y^*A_tN_t$$

In steady state, the growth rate of output equals the rate of population growth n plus the rate of technical progress g. By implication, the growth rate of output is independent of the saving rate.

How about capital and output per worker. Both grow at the rate of technical progress g, since, in steady state, capital and output per worker are given by,

$$K_t/N_t=k^*A_t$$

$$Y_t/N_t=y^*A_t$$

In steady state, the growth rate of output per worker equals the rate of technical progress g. By implication, the growth rate of output is independent of the saving rate. As in the model without technical progress and population growth, the savings rate determines output per effective unit of labor in the steady state.

The steady state is often referred to as the *balanced growth path*, in the sense that all aggregate and per capita variables (capital, output, consumption etc) grow at the same rate of exogenous growth determined by the sum of population growth and technical progress in the case of aggregates, or technical progress in the case of per capita aggregates.

Effects of a Rise in the Savings Rate on Capital and Output per Effective Unit of Labor

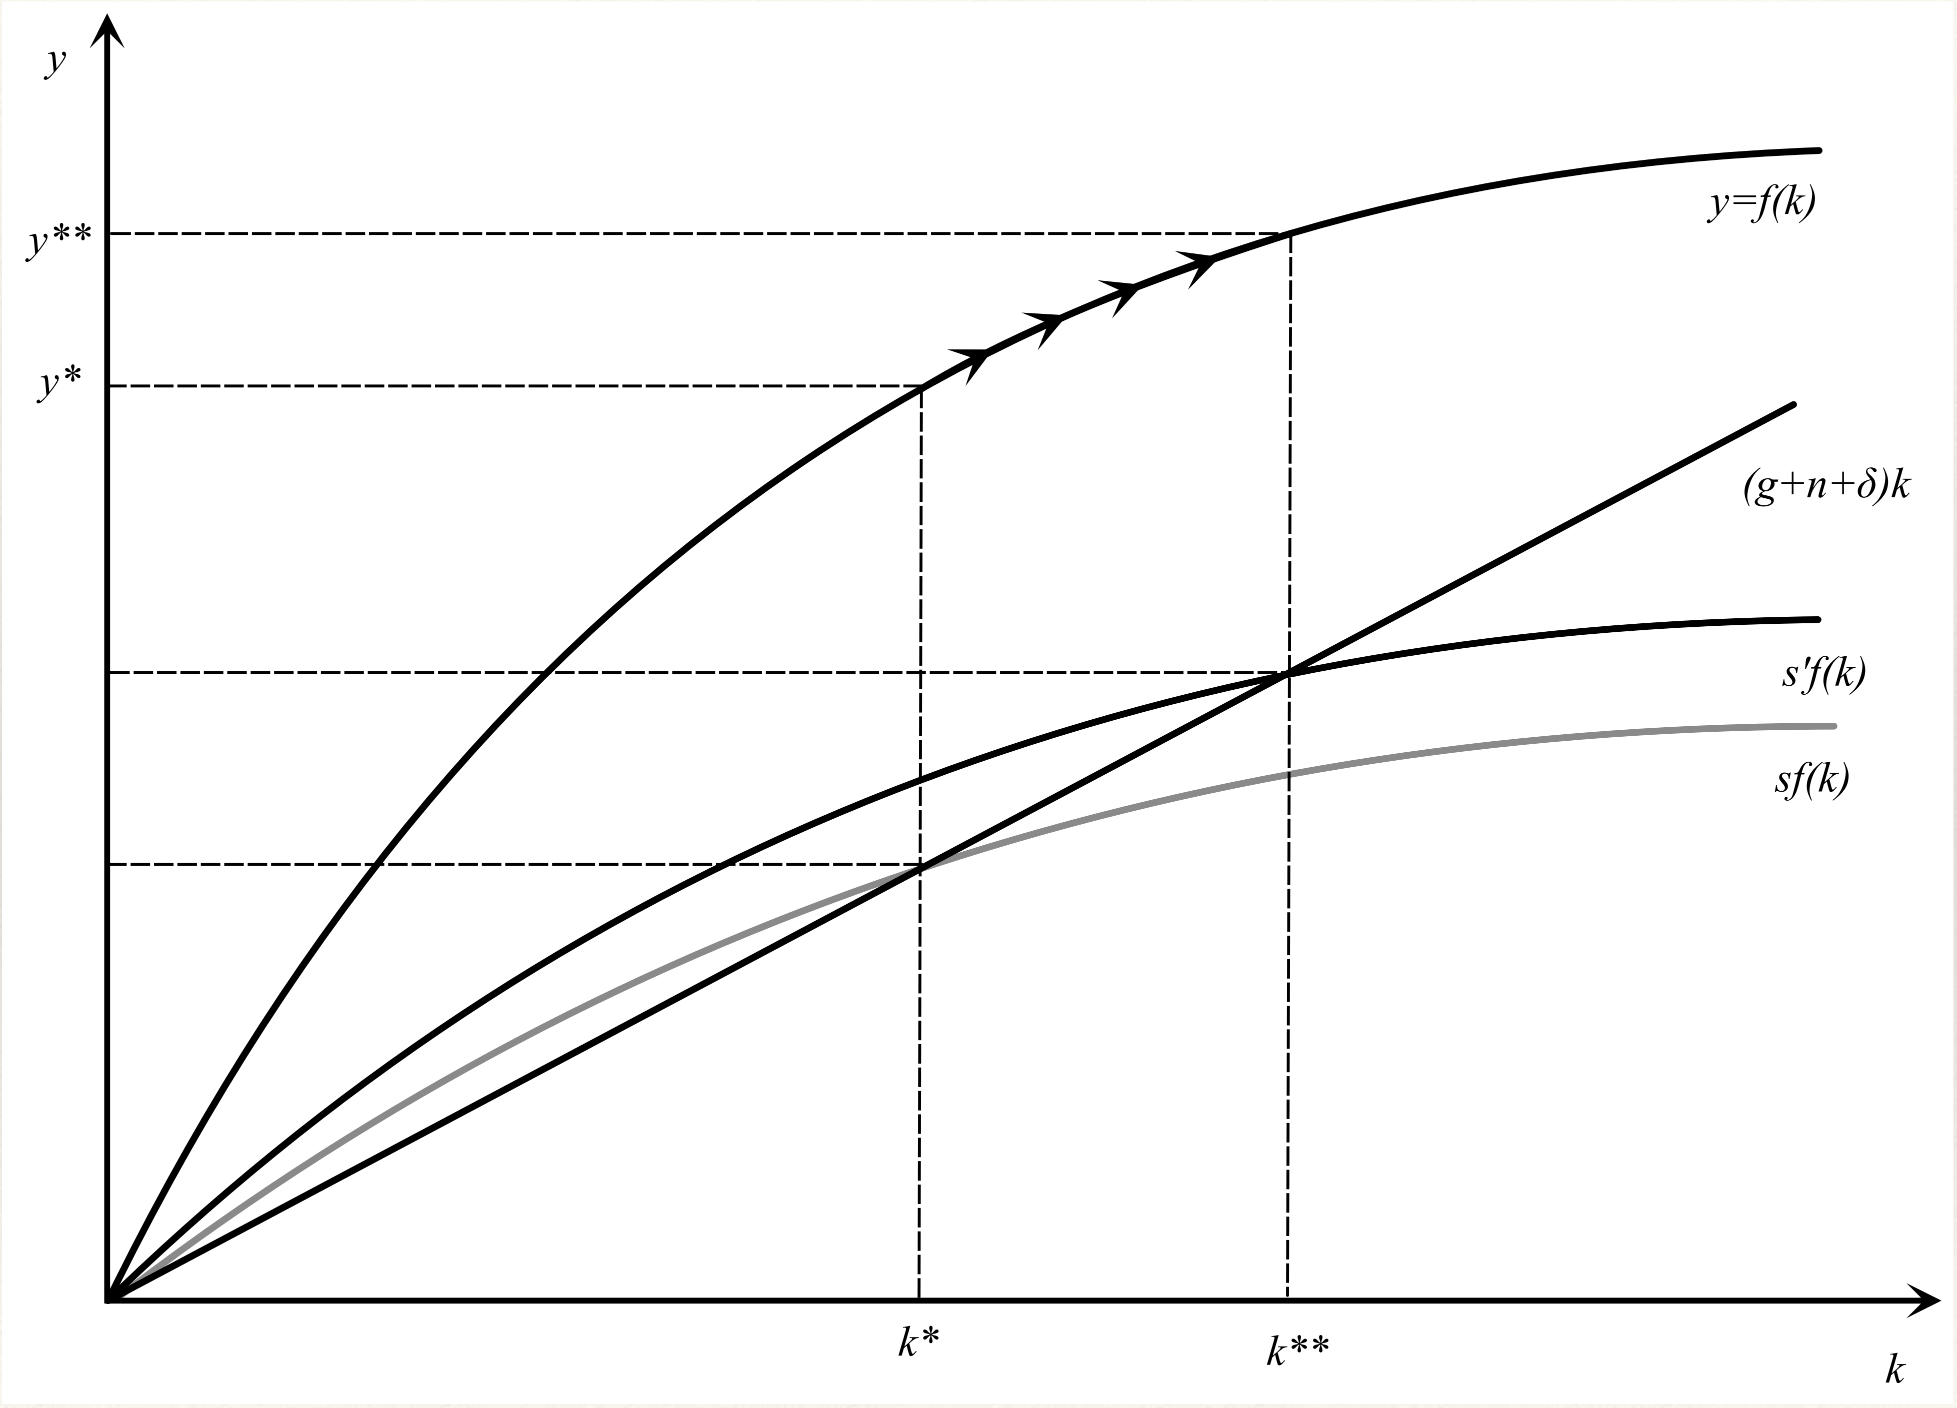

Effects of a Rise in the Savings Rate on Output per Worker

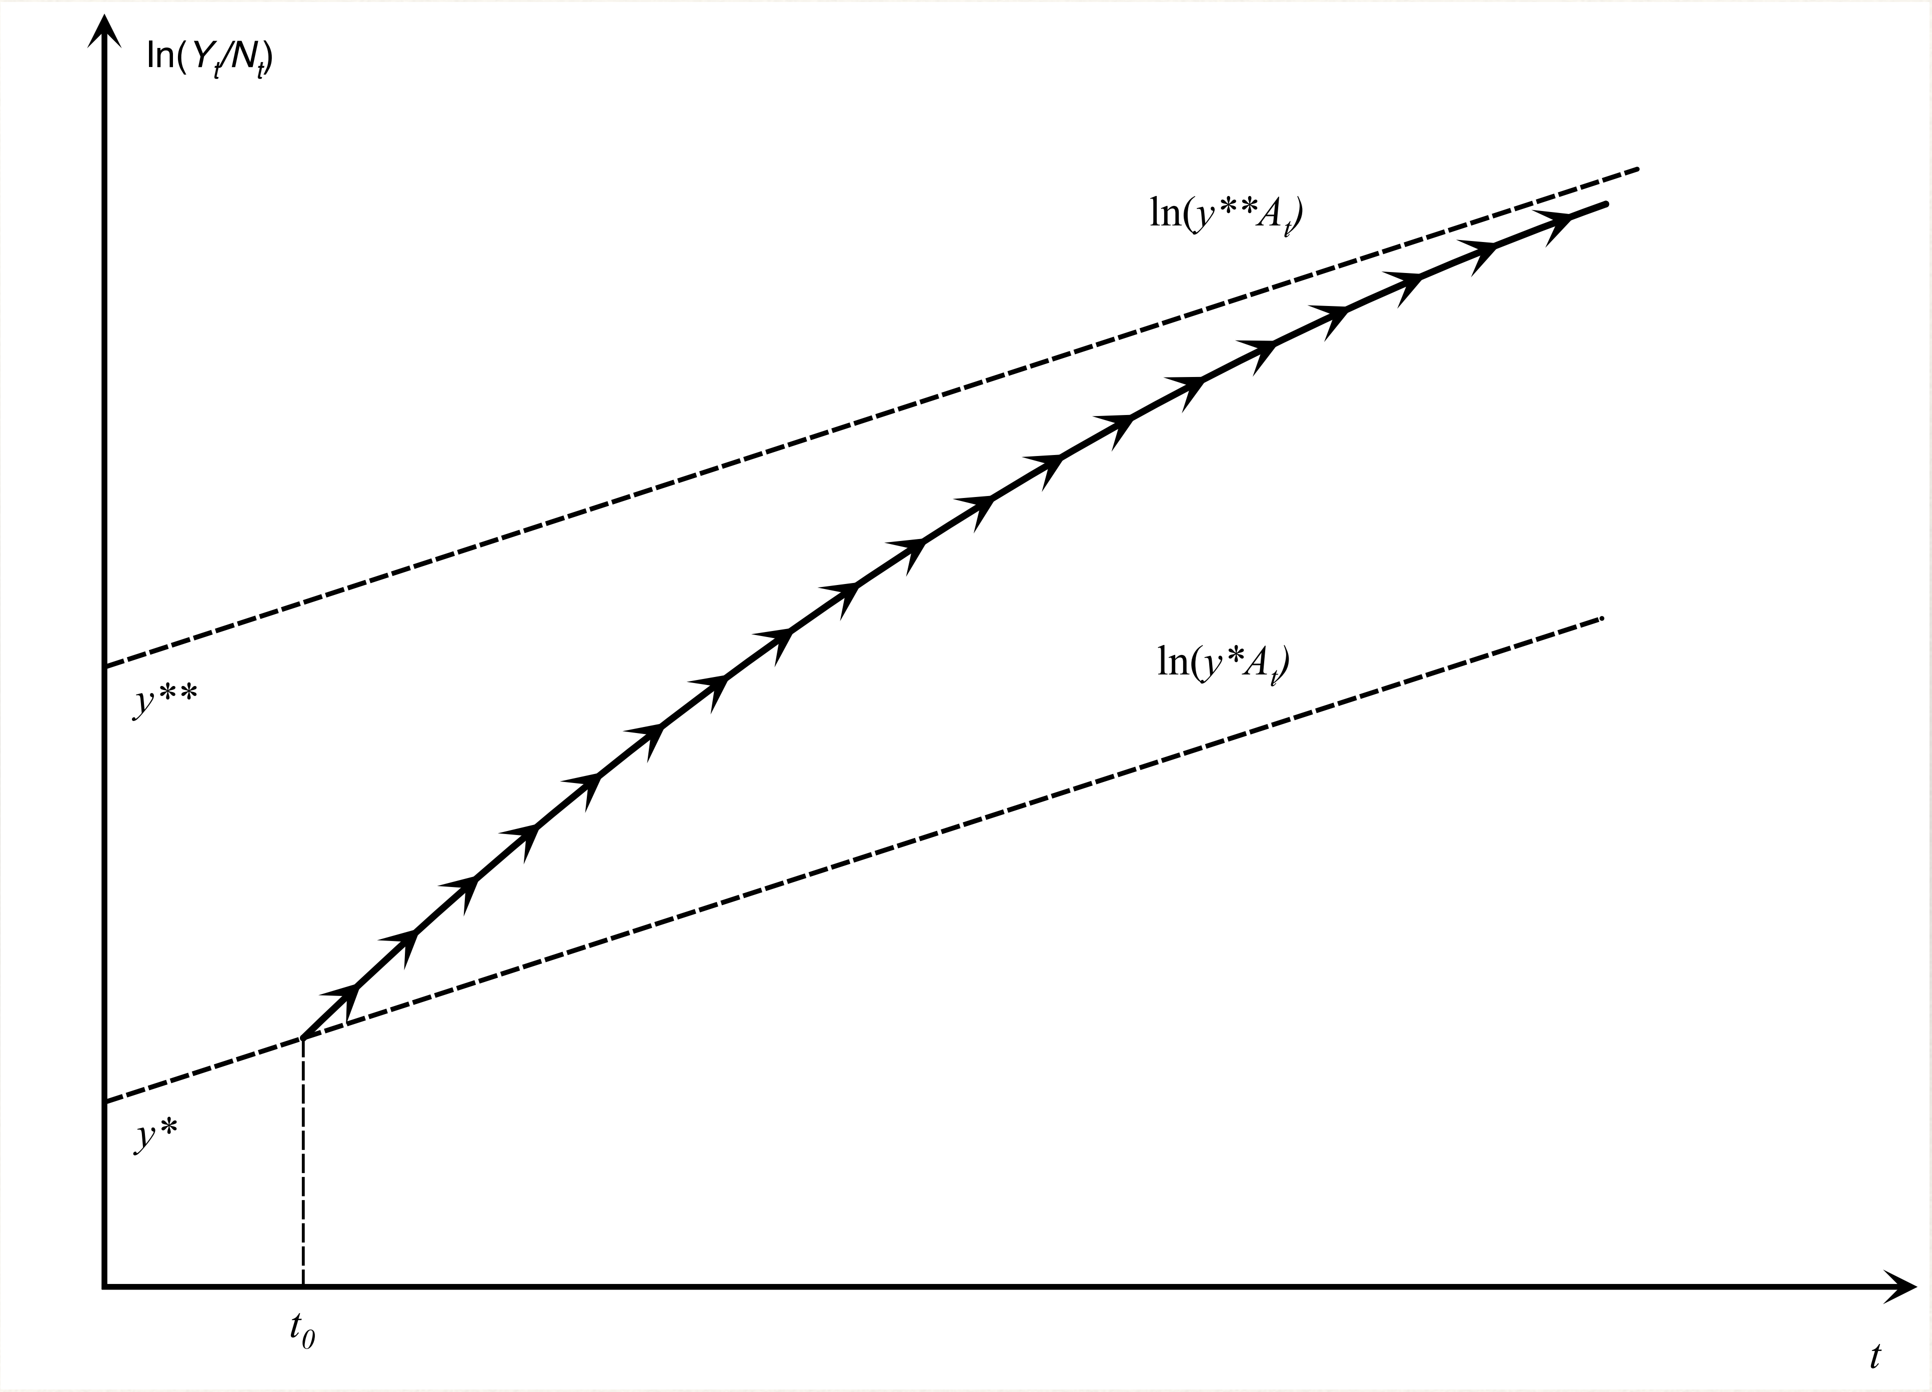

Conclusions about Economic Growth in the Presence of Technical Progress and Population Growth

1. In steady state, output per effective unit of labor and capital per effective unit of labor are constant. Thus, output per worker and capital per worker grow at the rate of technical progress.

2. In steady state, output and capital grow at the same rate as effective labor, and therefore at a rate equal to the growth rate of the number of workers plus the rate of technical progress. When the economy is in steady state, it is said to be on a balanced growth path.

3. The rate of output growth in steady state is independent of the saving rate. However, the saving rate affects the steady-state level of output per effective unit of labor and hence the steady state level of output per worker.

4. Finally, increases in the saving rate lead, for some time, to an increase in the growth rate above the steady-state growth rate.

What Determines Technical Progress

"Technical progress" brings to mind images of major discoveries: the invention of the microchip, the discovery of the structure of DNA, and so on. These discoveries suggest a process driven largely by scientific research and chance rather than by economic forces. But the truth is that most technical progress in modern economies is the result of a humdrum process: the outcome of firms' research and development (R&D) activities. Industrial R&D expenditures account for between 2% and 3% of GDP in each of the five major rich countries, the United States, Japan, Germany, France and the United Kingdom. About 75% of the roughly one million U.S. scientists and researchers working in R&D are employed by firms. U.S. firms' R&D spending equals more than 20% of their spending on gross investment, and more than 60% of their spending on net investment—gross investment less depreciation.

Firms spend on R&D for the same reason they buy new machines or build new plants: to increase profits. By increasing spending on R&D, a firm increases the probability that it will discover and develop a new "product". If the new product is successful, the firm's profits will increase.

There is, however, an important difference between purchasing a machine and spending more on R&D. The difference is that the outcome of R&D is fundamentally ideas. And, unlike a machine, an idea can potentially be used by many firms at the same time. A firm that has just acquired a new machine does not have to worry that another firm will use that particular machine. A firm that has discovered and developed a new product can make no such assumption. This last point implies that the level of R&D spending depends not only on the *fertility* of research - how spending on R&D translates into new ideas and new products - but also on the existence and the nature of *intellectual property rights* on research, that is the extent to which firms benefit from the results of their own R&D.

If firms cannot appropriate the profits from the development of new products, they will not engage in R&D and technical progress will be slow. One of the most important factors in this respect is the legal protection given to new products. Without such legal protection, profits from developing a new product are likely to be small. Except in rare cases where the product is based on a trade secret (such as Coca Cola), it will generally not take long for other firms to produce the same product, eliminating any advantage the innovating firm may have initially had. This is why countries have *patent laws*. Patents give a firm that has discovered a new product - usually a new technique or device - the right to exclude anyone else from the production or use of the new product for some time.

The Fertility of the Research Process

If research is very fertile—that is, if R&D spending leads to many new products—then, other things being equal, firms will have strong incentives to spend on R&D; R&D spending and, by implication, technical progress will be high. The determinants of the fertility of research lie largely outside the realm of economics. Many factors interact here:

The fertility of research depends on the successful interaction between *basic research* (the search for general principles and results) and *applied research and development* (the application of these results to specific uses, and the development of new products). Basic research does not lead, by itself, to technical progress. But the success of applied research and development depends ultimately on basic research. Much of the computer industry's development can be traced to a few breakthroughs, from the invention of the transistor to the invention of the microchip.

Some countries appear more successful at basic research; other countries are more successful at applied research and development. Studies point to differences in the education system as one of the reasons why. For example, it is often argued that the French higher education system, with its strong emphasis on abstract thinking, produces researchers who are better at basic research than at applied research and development. Studies also point to the importance of a "culture of entrepreneurship," in which a big part of technical progress comes from the ability of entrepreneurs to organize the successful development and marketing of new products—a dimension where the United States appears better than most other countries.

It takes many years, and often many decades, for the full potential of major discoveries to be realized. The usual sequence is one in which a major discovery leads to the exploration of potential applications, then to the development of new products, and, finally, to the adoption of these new products. Thirty five years after the commercial introduction of personal computers, it often seems as if we have just begun discovering their uses.

An age-old worry is that research will become less and less fertile, that most major discoveries have already taken place and that technical progress will begin to slow down. This fear may come from thinking about mining, where higher-grade mines were exploited first, and where we have had to exploit increasingly lower-grade mines. But this is only an analogy, and so far there is no evidence that it is correct.

The Protection of Intellectual Property Rights

How should governments design patent laws? On the one hand, protection is needed to provide firms with the incentives to spend on R&D. On the other, once firms have discovered new products, it would be best for society if the knowledge embodied in those new products were made available to other firms and to people without restrictions.

Take, for example, biogenetic research. Only the prospect of large profits is leading bio-engineering firms to embark on expensive research projects. Once a firm has found a new product, and the product can save many lives, it would clearly be best to make it available at cost to all potential users. But if such a policy was systematically followed, it would eliminate incentives for firms to do research in the first place. So, patent law must strike a difficult balance. Too little protection will lead to little R&D. Too much protection will make it difficult for new R&D to build on the results of past R&D, and may also lead to little R&D. This type of dilemma is known as "time inconsistency" and it applies in other areas too, such as the taxation of capital and monetary policy.

Countries that are less technologically advanced often have poorer patent protection. Emerging economies, such as China, for example, are usually associated with poor enforcement of intellectual property and patent rights. Our discussion helps explain why. These countries are typically users rather than producers of new technologies. Much of their improvement in productivity comes not from inventions within the country, but from the adaptation of foreign technologies. In this case, the costs of weak patent protection are small, because there would be few domestic inventions anyway. But the benefits of low patent protection are clear: They allow domestic firms to use and adapt foreign technology without having to pay royalties to the foreign firms that developed the technology, which is good for the country itself, but not for technical progress in general.

Revisiting the Facts of Growth

Suppose we observe an economy with a high growth rate of output per worker over some period of time. Our theory implies this fast growth may come from two sources:

A. It may reflect a high rate of technical progress under balanced growth.

B. It may reflect convergence, i.e the adjustment of capital per effective worker, to a higher level.

Can we tell how much of the growth comes from one source and how much comes from the other? Yes. If high growth reflects high balanced growth, output per worker should be growing at a rate equal to the rate of technical progress. If high growth reflects instead the adjustment to a higher level of capital per effective worker, this adjustment should be reflected in a growth rate of output per worker that exceeds the rate of technical progress.

Looking at the experience of the major developed economies since 1985, it appears as if growth has come from technical progress, not unusually high capital accumulation. This does not say that capital accumulation was irrelevant. Capital accumulation was such as to allow these countries to maintain a roughly constant ratio of output to capital and achieve balanced growth.

A second conclusion is that convergence of output per worker between the United States and the other industrial countries comes from higher technological progress rather than from faster capital accumulation. All the other industrial countries started behind the United States in 1985. In all these countries the rate of technical progress has been higher than in the United States.

This is an important conclusion. One can think, in general, of two sources of convergence between countries. First: Poorer countries are poorer because they have less capital to begin with. Over time, they accumulate capital faster than the others, generating convergence. Second: Poorer countries are poorer because they are less technologically advanced than the others. Thus, over time, they become more sophisticated, either by importing technology from advanced countries or developing their own. As technological levels converge, so does output per worker. The conclusion we can draw is that, in the case of rich countries, the more important source of convergence since 1985 has been technological convergence. On the other hand, convergence between the emerging economies and the industrial economies is due to both capital accumulation and technological convergence.

Technical Progress in Advanced and Emerging Economies

The nature of technical progress is likely to be different between more advanced and less advanced economies.

The more advanced economies, being by definition at the technological frontier, need to develop new ideas, new processes, new products. They need to innovate.

The countries that are behind can instead improve their level of technology by copying and adapting the new processes and products developed in the more advanced economies. They need to imitate. The farther behind a country is, the larger the role of imitation relative to innovation.

As imitation is likely to be easier than innovation, this can explain why convergence, both within the OECD, and in the case of China and other emerging economies, typically takes the form of technological catch-up.

It raises, however, yet another question: If imitation is so easy, why is it that so many other countries do not seem to be able to do the same and grow? This points to the broader aspects of technology we discussed earlier in the chapter. Technology is more than just a set of blueprints. How efficiently these blueprints can be used and how productive an economy is depend on its institutions, on the quality of its government, and so on.