The steady-state and transition dynamics are obtained along reasoning analogous to the one employed above. In equilibrium, income per efficiency unit remains constant. Since *efficiency units* of labour grow faster than labour, due to technological progress, output (and capital) per worker must be growing. To show this mathematically, we may start by noting that in the steady state income per efficiency units of labour does not change, $\Delta\hat{y} = 0$. From the definition $\hat{y} \equiv \frac{Y}{EL}$ we obtain per capita income y by multiplying by E:

$$y \equiv \frac{Y}{L} = \frac{Y}{EL}E = \hat{y} \times E$$

Finally, we recall that the growth rate of the product $\hat{y} \times E$ can be approximated by the sum of the growth rates of $\hat{y}$ and E:

$$\frac{\Delta y}{y} = \frac{\Delta\hat{y}}{\hat{y}} + \frac{\Delta E}{E} = \frac{\Delta\hat{y}}{\hat{y}} + \varepsilon = 0 + \varepsilon = \varepsilon$$

This shows that even though income per efficiency units of labour does not change in the steady state, $\Delta\hat{y} = 0$, income per capita nevertheless does. It grows at the rate of technological progress ε. So we finally have a model that explains *income growth* in the conventional meaning of the term.

As regards comparative statics, a faster rate of technological progress turns the requirement curve upwards, thus lowering capital and income *per efficiency unit*. Does this mean that faster technological progress is bad? With regard to per capita income, the answer is no. Remember that the one-off technology improvement analyzed in section 9.4 raised capital and output *per worker*. The same result must apply here, where the one-off technological improvement simply occurs period after period. Therefore, *faster technological progress raises the level and the growth rate of output per worker*.

CASE STUDY 9.2 Income and leisure choices in the OECD countries

When microeconomists analyze individual behaviour they usually assume that two things enhance a person's utility: first, consumption (which is limited by income); second, leisure time (the time we have to enjoy the things we consume). This makes it obvious that judging the well-being of a country's citizens by looking at income would be just as one-sided as judging their well-being by looking at leisure time.

Using data for the year 1996, Figures 1 and 2 show that a country's per capita income and its leisure time need not necessarily go hand in hand. Figure 1 shows per capita incomes relative to the OECD average normalized to 100. The richest country in the sample is the USA, with per capita income 35% above average. The poorest country is Portugal, whose income falls short of the OECD

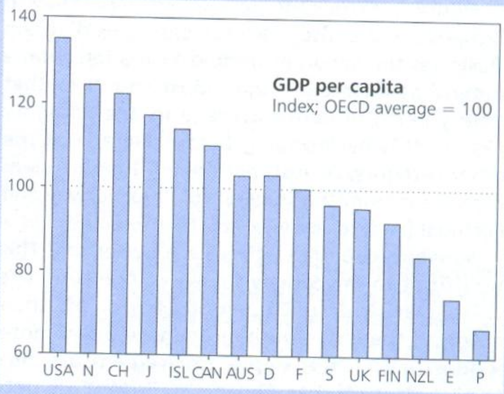

Figure 1

3 Note that the vertical distance between the production function and the requirement line measures consumption available at different steady states.
4 Consumption is maximized where a line parallel to the requirement line just touches the production function. This point defines golden-rule output and the golden-rule capital stock.
5 Since the actual savings curve must intersect the requirement line at the golden-rule capital stock, this identifies the golden-rule savings rate.

Dynamic efficiency

If the actual savings rate does not correspond with the savings rate recommended by the golden rule, should the government try to move it towards s_{gold}, say by offering tax incentives? Well, that depends.

Assume first that the savings rate is too high, and that this led to the steady-state capital stock K_1^* and a level of consumption C_1^* that falls short of maximum steady-state consumption C_{gold}^* (see Figure 9.14). When citizens change their behaviour, lowering the savings rate from s_1 to s_{gold}, consumption rises immediately to C_1'. Subsequently, consumption gradually falls as the capital stock begins to melt away, but it will always remain higher than C_1^*. The time path of consumption looks as displayed in the left panel of Figure 9.15. To reduce the savings rate from s_1 to s_{gold} would provide individuals

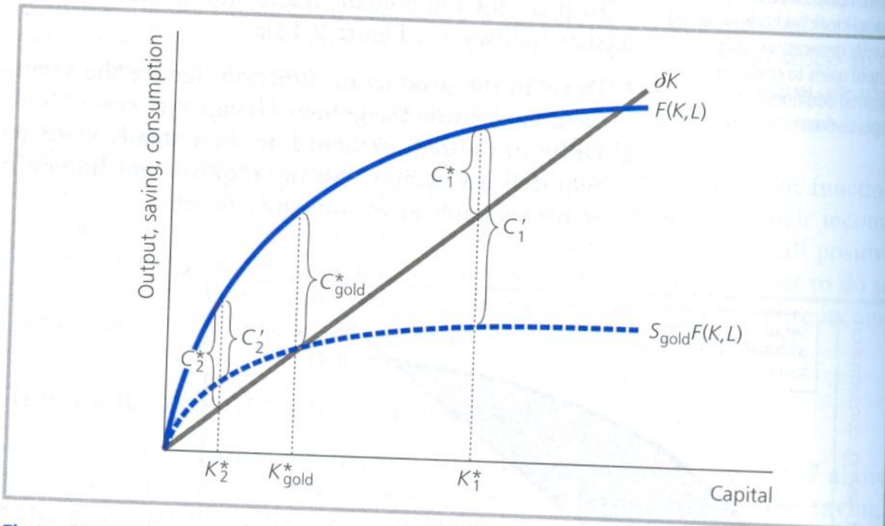

Figure 9.14 When the savings rate exceeds s_{gold}, a steady state capital stock such as K_1^* results, and consumption is C_1^*. When lowering the savings rate to s_{gold}, the immediate effect on consumption is a drop to C_1'. While the capital stock subsequently shrinks towards K_{gold}^*, consumption is always given by the vertical distance between the production function and the savings function. It exeeds C_1^* at all points in time. When the savings rate falls short of s_{gold}, a steady state capital stock such as K_2^* results, and consumption is C_2^*. After raising the savings rate to s_{gold}, consumption initially falls to C_2'. While the higher savings rate makes the capital stock grow towards K_{gold}^*, consumption remains as given by the vertical distance between the production function and the savings function. It is initially smaller than C_2^*, but later surpasses it and remains higher for good.

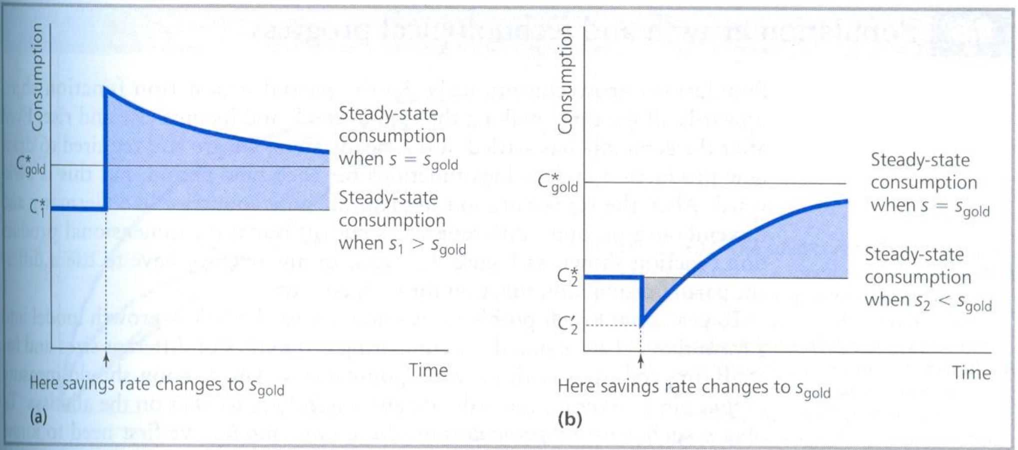

Figure 9.15 Savings rates smaller or larger than that required by the golden rule restrict the country to lower steady-state consumption. Paths of adjustment from the old, suboptimal steady state to the new, golden-steady state differ in the two cases shown in Figure 9.15, panels (a) and (b). If $s > s_{gold}$, reducing the savings rate to s_{gold} improves consumption now and forever (panel (a)). The country would gain all the consumption indicated by the area tinted blue if it adopted s_{gold}. Sticking to s_1 is dynamically inefficient. If $s < s_{gold}$, the country faces a dilemma (panel (b)). Raising s to s_{gold} only pays off later in the form of consumption gains tinted blue. Before consumption improves, the country goes through a period of reduced consumption. These losses are tinted grey.

with higher consumption today and during all future periods – at no cost. The sum of all consumption gains, compared to the initial steady state, is represented by the area shaded blue. Not to jump at the opportunity to reap this costless gain would be foolish or irrational – or inefficient. This is why a steady state like K^*_1, or any other steady-state capital stock that exceeds the golden one, is called **dynamically inefficient.**

Things are different when the savings rate is too low, say, at s_2. Then the steady-state capital stock K^*_2 obtains, and, again, the accompanying level of consumption C^*_2 falls short of C^*_{gold} (Figure 9.14). To put the economy on a path towards the golden steady state, the savings rate needs to increase from s_2 to s_{gold}. While this will succeed in raising consumption in the long run, the price to pay is an immediate drop in consumption from C^*_2 to C'_2. Only as the higher savings rate leads to capital accumulation and growing income does consumption recover and, at some point in time, surpass its initial level (Figure 9.15, panel (b)). Consumption in the more distant future can only be raised at the cost of reduced consumption in the short and medium run. The consumption loss incurred in the early periods (shaded grey) is the price for the longer-run consumption gains (shaded blue). So the question boils down to how much weight we want to put on today's (or *this* generation's) consumption as compared to tomorrow's (or *future* generation's) consumption. This is not for the economist to decide. His or her proper task is to set out the options. But when future benefits are being discounted heavily compared to current costs, it is not necessarily irrational not to raise the savings rate from s_2 to s_{gold}. This is why a steady state like K^*_2, or any other steady-state capital stock that falls short of the golden one, is called **dynamically efficient.**

Empirical note. Most countries save less and, hence, accumulate less capital than the golden rule suggests. Thus, they do face the dilemma of whether to reduce today's consumption in order to raise tomorrow's.

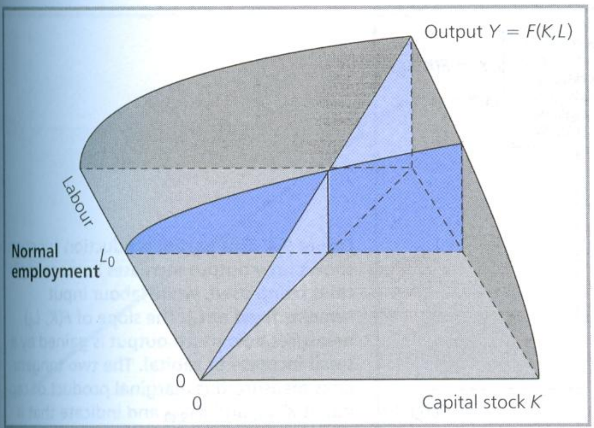

Figure 9.3 The 3D production function shows how, for a given production technology, output rises as greater and greater quantities of capital and/or labour are being employed. As a reminder, for first and second derivatives we assume $F_K, F_L > 0$ and $F_{KK}, F_{LL} < 0$.

Figure 9.3 displays this function again, which is called the **extensive form** of the production function. Note, however, that the axes have been relabelled. This is because we now shift our perspective. In Chapter 6, when deriving the labour demand curve, we asked how at any point in time, with a given capital stock that could not be changed in the short run, different amounts of labour employed by firms would affect output produced.

Here we want to know why a country has the capital stock it has. To obtain an unimpaired view on this issue, we now ignore the business cycle. For a start we assume that employment is fixed at normal employment L_0, at which the labour market clears. In order not to have to differentiate all the time between magnitudes per capita or per worker, we even suppose that all people work. So the number of workers equals the population. All our arguments go through, however, if workers are a fixed share of the population. If this share changes, the effects are analogous to what results from a changing population as will be discussed in section 9.6.

The assumptions that economists make about the production function shown in Figure 9.3 are (adding a third one) as follows:

- Output increases as either factor or both factors increase.
- If one factor remains fixed, increases of the other factor yield smaller and smaller output gains.
- If both factors rise by the same percentage, output also rises by this percentage.

As we know from Chapter 6, the second assumption refers to partial production functions. For our current purposes we place a vertical cut through the production function parallel to the axis measuring the capital stock. Figure 9.4 shows the obtained partial production function that fixes labour at L_0.

What we said about the partial production function employed in Chapter 6 applies in a similar way to the one displayed in Figure 9.4. The output gain accomplished by a small increase in K (which is called the **marginal product of capital**) is measured by the slope of the production

The **marginal product of capital** is the output added by adding one unit of capital.

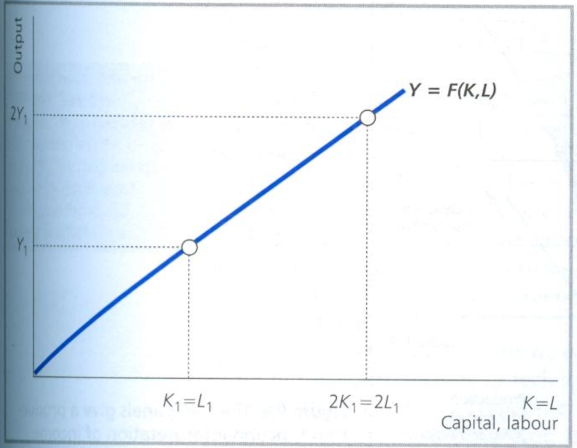

Figure 9.5 This production function shows how output increases as capital and labour rise in proportion. $F(K, L = K)$ is a straight line, indicating that we assume *constant returns to scale*: if capital and labour increase by a given percentage, output increases by the same percentage.

factors that enter the production function, without asking why those factors developed the way they did. This question is left to **growth theory**, to which we will turn below.

Note. The formulation of this particular functional form as a basis for empirical estimates is due to US economist turned politician Paul Douglas and mathematician Charles Cobb.

As the word 'accounting' implies, growth accounting wants to arrive at some hard numbers. A general function like equation (9.1) is not useful for this purpose. Economists therefore use more specific functional forms when turning to empirical work. The most frequently employed form is the **Cobb–Douglas production function:**

$$Y = AK^{\alpha}L^{1-\alpha} \qquad \textbf{Cobb–Douglas production function} \qquad (9.2)$$

As Box 9.1 shows, this function has the same properties assumed to hold for the general production function discussed above, plus a few other properties that come in handy during mathematical operations and appear to fit the data quite well.

Equation (9.2) states that income is related to the factor inputs K and L and to the production technology as measured by the leading variable A. This leaves two ways for economic growth to occur, as Figure 9.6 illustrates. In panel (a) we keep technology constant between 1950 and the year 2000. Income grows only because of an expanding capital stock and a growing labour force. In panel (b) technology has improved, tilting the production function upwards. As a consequence GDP rises at any given combination of capital and labour employed.

The two motors of economic growth featured in the two panels of Figure 9.6 operate simultaneously. Growth accounting tries to identify their qualitative contributions. This is tricky, since the three factors comprising the multiplicative term on the right-hand side of equation (9.2) interact, affecting each other's contribution. A first step towards disentangling this is to take natural logarithms. This yields

$$\ln Y = \ln A + \alpha \ln K + (1 - \alpha)\ln L \qquad (9.3)$$

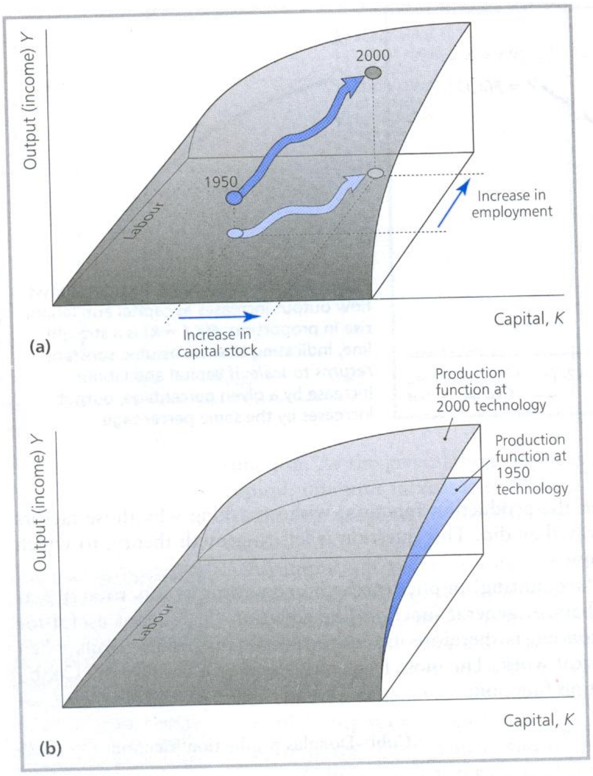

Figure 9.6 The two panels give a production function interpretation of income growth. Panel (a) assumes constant production technology. Then the production function graph does not change in this diagram. Income has nevertheless grown from 1950 to 2000 because the capital stock has risen and employment has gone up. Panel (b) illustrates the effect of technological progress on the production function graph. The upwards tilt of the production function would raise income even if input factors did not change. In reality all three indicated causes of income growth play a role: capital accumulation, labour force growth and technological progress.

Source: K. Case, R. Fair, M. Gärtner and K. Heather (1999), *Economics*, Harlow: Prentice Hall Europe.

Maths note. An alternative way to derive the growth-accounting equation starts by taking the total differential of the production function $Y = AK^{\alpha}L^{1-\alpha}$ which is $dY = K^{\alpha}L^{1-\alpha}dA + \alpha AK^{\alpha-1}L^{1-\alpha}\,dK + (1-\alpha)AK^{\alpha}L^{-\alpha}dL$. Now divide by Y on the left-hand side and by $AK^{\alpha}L^{1-\alpha}$ on the right-hand side to obtain (after cancelling terms) $\frac{dY}{Y} = \frac{dA}{A} + \alpha\frac{dK}{K} + (1-\alpha)\frac{dL}{L}$ which is the continuous-time analogue to equation (9.4).

meaning that the logarithm of income is a weighted sum of the logarithms of technology, capital and labour. Now take first differences on both sides (meaning that we deduct last period's values) to obtain $\ln Y - \ln Y_{-1} = \ln A - \ln A_{-1} + \alpha(\ln K - \ln K_{-1}) + (1-\alpha)(\ln L - \ln L_{-1})$. Finally, making use of the property (mentioned previously and derived in the appendix on logarithms in Chapter 1) that the first difference in the logarithm of a variable is a good approximation for this variable's growth rate, we arrive at

$$\frac{\Delta Y}{Y} = \frac{\Delta A}{A} + \alpha\frac{\Delta K}{K} + (1-\alpha)\frac{\Delta L}{L} \qquad \text{Growth accounting equation} \qquad (9.4)$$

stating that a country's income growth is a weighted sum of the rate of technological progress $\Delta A/A$, capital growth and employment growth. All we need to know now before we can do some calculations with this equation is the magnitude of α. This is not as hard as it may seem, at least not if we assume that our economy operates under perfect competition. Perfect competition ensures that each factor of production is paid the marginal product it generates. As we already saw in Chapter 6 in the context of the