CHAPTER 1

Introduction

Paul Samuelson once stated that "macroeconomics, even with all of our computers and with all of our information is not an exact science and is incapable of being an exact science". Perhaps this quote captures the view that the field of macroeconomics, the study of aggregate behaviour of the economy, is full of loose ends and inconsistent statements that make it difficult for economists to agree on anything.

While there is truth to the fact that there are plenty of disagreements among macroeconomists, we believe such a negative view is unwarranted. Since the birth of macroeconomics as a discipline in the 1930s, in spite of all the uncertainties, inconsistencies, and crises, macroeconomic performance around the world has been strong. More recently, dramatic shocks, such as the Great Financial Crisis or the Covid pandemic, have been managed – not without cost, but with effective damage control. There is much to celebrate in the field of macroeconomics.

Macroeconomics was born under the pain of both U.S. and UK's protracted recession of the 1930s. Until then, economics had dealt with markets, efficiency, trade, and incentives, but it was never thought that there was place for a large and systematic breakdown of markets. High and persistent unemployment in the U.S. required a different approach.

The main distinctive feature to be explained was the large disequilibrium in the labour market. How could it be that a massive number of people wanted to work, but could not find a job? This led to the idea of the possibility of aggregate demand shortfalls – and thus of the potential role for government to prop it up, and, in doing so, restore economic normalcy. "Have people dig a hole and fill them up if necessary" is the oft-quoted phrase by Keynes. In modern economic jargon, increase aggregate demand to move the equilibrium of the economy to a higher level of output.

Thus, an active approach to fiscal and monetary policy developed, entrusting policy makers with the role of moderating the business cycle. The relationship was enshrined in the so-called Phillips curve, a relationship that suggested a stable tradeoff between output and inflation. If so, governments simply had to choose their preferred spot on that tradeoff.

Then things changed. Higher inflation in the 60s and 70s, challenged the view of a stable tradeoff between output and inflation. In fact, inflation increased with no gain in output, the age of stagflation had arrived. What had changed?

The answer had to do with the role of expectations in macroeconomics.[1]

The stable relationship between output and inflation required static expectations. People did not expect inflation, then the government found it was in its interest to generate a bit of inflation – but

How to cite this book chapter:
Campante, F., Sturzenegger, F. and Velasco, A. 2021. *Advanced Macroeconomics: An Easy Guide.* Ch. 1. 'Introduction', pp. 1–4. London: LSE Press. DOI: https://doi.org/10.31389/lsepress.ame.a

that meant people were always wrong! As they started anticipating the inflation, then its effect on employment faded away, and the effectiveness of macro policy had gone stale.

The rational expectations revolution in macroeconomics, initiated in the 1970s, imposed the constraint that a good macro model should allow agents in the model to understand it and act accordingly. This was not only a theoretical purism. It was needed to explain what was actually happening in the real world. The methodological change took hold very quickly and was embraced by the profession. As a working assumption, it is a ubiquitous feature of macroeconomics up to today.

Then an additional challenge to the world of active macroeconomic policy came about. In the early 1980s, some macroeconomists started the "real business cycles" approach: they studied the neoclassical growth model – that is, a model of optimal capital accumulation – but added to it occasional productivity shocks. The result was a simulated economy that, they argued, resembled on many dimensions the movements of the business cycle. This was a dramatic finding because it suggested that business cycles could actually be the result of optimal responses by rational economic agents, thereby eschewing the need for a stabilising policy response. What is more, active fiscal or monetary policy were not merely ineffective, as initially argued by the rational expectations view: they could actually be harmful.

This was the state of the discussion when a group of economists tackled the task of building a framework that recovered some of the features of the old Keynesian activism, but in a model with fully rational agents. They modelled price formation and introduced market structures that departed from a perfectly competitive allocation. They adhered strictly to the assumptions of rational expectations and optimisation, which had the added advantage of allowing for explicit welfare analyses. Thus, the New Keynesian approach was built. It also allowed for shocks, of course, and evolved into what is now known as dynamic stochastic general equilibrium (DSGE) models.

Macroeconomic policymaking evolved along those lines. Nowadays, DSGE models are used by any respectable central bank. Furthermore, because this type of model provides flexibility in the degree of price rigidities and market imperfections, it comprises a comprehensive framework nesting the different views about how individual markets operate, going all the way from the real business cycle approach to specifications with ample rigidities.

But the bottom line is that macroeconomics speaks with a common language. While differences in world views and policy preferences remain, having a common framework is a great achievement. It allows discussions to be framed around the parameters of a model (and whether they match the empirical evidence) – and such discussions can be more productive than those that swirl around the philosophical underpinnings of one's policy orientations.

This book, to a large extent, follows this script, covering the different views – and very importantly, the tools needed to speak the language of modern macroeconomic policymaking – in what we believe is an accessible manner. That language is that of dynamic policy problems.

We start with the Neoclassical Growth Model – a framework to think about capital accumulation through the lens of optimal consumption choices – which constitutes the basic grammar of that language of modern macroeconomics. It also allows us to spend the first half of the book studying economic growth – arguably the most important issue in macroeconomics, and one that, in recent decades, has taken up as much attention as the topic of business cycles. The study of growth will take us through the discussion of factor accumulation, productivity growth, the optimality of both the capital stock and the growth rate, and empirical work in trying to understand the proximate and fundamental causes of growth. In that process, we also develop a second canonical model in modern macroeconomics: the overlapping generations model. This lets us revisit some of the issues around capital accumulation and long-run growth, as well as study key policy issues, such as the design of pension systems.

We then move to discuss issues of consumption and investment. These are the key macroeconomic aggregates, of course, and their study allows us to explore the power of the dynamic tools we developed in the first part of the book. They also let us introduce the role of uncertainty and expectations, as well as the connections between macroeconomics and finance.

Then, in the second half of the book, we turn to the study of business cycle fluctuations, and what policy can and should do about it. We start with the real business cycle approach, as it is based on the neoclassical growth model. Then we turn to the Keynesian approach, starting from the basic IS-LM model, familiar to anyone with an undergraduate exposure to macroeconomics, but then showing how its modern version emerged: first, with the challenge of incorporating rational expectations, and then with the fundamentals of the New Keynesian approach. Only then, we present the canonical New Keynesian framework.

Once we've covered all this material, we discuss the scope and effectiveness of fiscal policy. We also go over what optimal fiscal policy would look like, as well as some of the reasons for why in practice it departs from those prescriptions. We then move to discuss monetary policy: the relationship between money and prices, the debate on rules vs discretion, and the consensus that arose prior to the 2008 financial crisis and the Great Recession. We then cover the post-crisis development of quantitative easing, as well as the constraints imposed by the zero lower bound on nominal interest rates. We finish off by discussing some current topics that have been influencing the thinking of policymakers on the fiscal and monetary dimensions: secular stagnation, the fiscal theory of the price level, and the role of asset-price bubbles and how policy should deal with them.

As you can see from this whirlwind tour, the book covers a lot of material. Yet, it has a clear methodological structure. We develop the basic tools in the first part of the book, making clear exactly what we need at each step. All you need is a basic knowledge of calculus, differential equations, and some linear algebra – and you can consult the mathematical appendix for the basics on the tools we introduce and use in the book. Throughout, we make sure to introduce the tools not for their own sake, but in the context of studying policy-relevant issues and helping develop a framework for thinking about dynamic policy problems. We then study a range of policy issues, using those tools to bring you to the forefront of macroeconomic policy discussions. At the very end, you will also find two appendices for those interested in tackling the challenge of running and simulating their own macroeconomic models.

All in all, Samuelson was right that macroeconomics cannot be an exact science. Still, there is a heck of a lot to learn, enjoy and discover – and this, we hope, will help you become an informed participant in exciting macroeconomic policy debates. Enjoy!

Note

[1] Surprisingly, the answer came from the most unexpected quarter: the study of agricultural markets. As early as 1960 John Muth was studying the cobweb model, a standard model in agricultural economics. In this model the farmers look at the harvest price to decide how much they plant, but then this provides a supply the following year which is inconsistent with this price. For example a bad harvest implies a high price, a high price implies lots of planting, a big harvest next year and thus a low price! The low price motivates less planting, but then the small harvest leads to a high price the following year! In this model, farmers were systematically wrong, and kept being wrong all the time. This is nonsense, argued Muth. Not only should they learn, they know the market and they should plant the equilibrium price, namely the price that induces the amount of planting that implies that next year that will be the price. There are no cycles, no mistakes, the market equilibrium holds from day one! Transferred to macroeconomic policy, something similar was happening.

Growth Theory

CHAPTER 2

Growth theory preliminaries

2.1 | Why do we care about growth?

It is hard to put it better than Nobel laureate Robert Lucas did as he mused on the importance of the study of economic growth for macroeconomists and for anyone interested in economic development.[1]

> 'The diversity across countries in measured per capita income levels is literally too great to be believed. (...) Rates of growth of real per capita GNP are also diverse, even over sustained periods. For 1960–80 we observe, for example: India, 1.4% per year; Egypt, 3.4%; South Korea, 7.0%; Japan, 7.1%; the United States, 2.3%; the industrial economies averaged 3.6%. (..) An Indian will, on average, be twice as well off as his grandfather; a Korean 32 times. (...) I do not see how one can look at figures like these without seeing them as representing possibilities. Is there some action a government of India could take that would lead the Indian economy to grow like Indonesia's or Egypt's? If so, what, exactly? If not, what is it about the 'nature of India' that makes it so? *The consequences for human welfare involved in questions like these are simply staggering: Once one starts to think about them, it is hard to think about anything else.*'
>
> Lucas Jr. (1988) (emphasis added)

While it is common to think about growth today as being somehow natural, even expected – in fact, if world growth falls from 3.5 to 3.2%, it is perceived as a big crisis – it is worthwhile to acknowledge that this was not always the case. Pretty much until the end of the 18th century growth was quite low, if it happened at all. In fact, it was so low that people could not see it during their lifetimes. They lived in the same world as their parents and grandparents. For many years it seemed that growth was actually *behind* as people contemplated the feats of antiquity without understanding how they could have been accomplished. Then, towards the turn of the 18th century, as shown in Figure 2.1 something happened that created explosive economic growth as the world had never seen before. Understanding this transition will be the purpose of Chapter 10. Since then, growth has become the norm. This is the reason the first half of this book, in fact up to Chapter 10, will deal with understanding growth. As we proceed we will ask about the determinants of capital accumulation (Chapters 2 through 5, as well as 8 and 9), and discuss the process of technological progress (Chapter 6). Institutional factors will be addressed in Chapter 7. The growth process raises many interesting questions: should we expect this growth to continue? Should we expect it eventually to decelerate? Or, on the contrary, will it accelerate without bound?

How to cite this book chapter:
Campante, F., Sturzenegger, F. and Velasco, A. 2021. *Advanced Macroeconomics: An Easy Guide.* Ch. 2. 'Growth theory preliminaries', pp. 7–22. London: LSE Press.
DOI: https://doi.org/10.31389/lsepress.ame.b

Figure 2.1 The evolution of the world GDP per capita over the years 1–2008

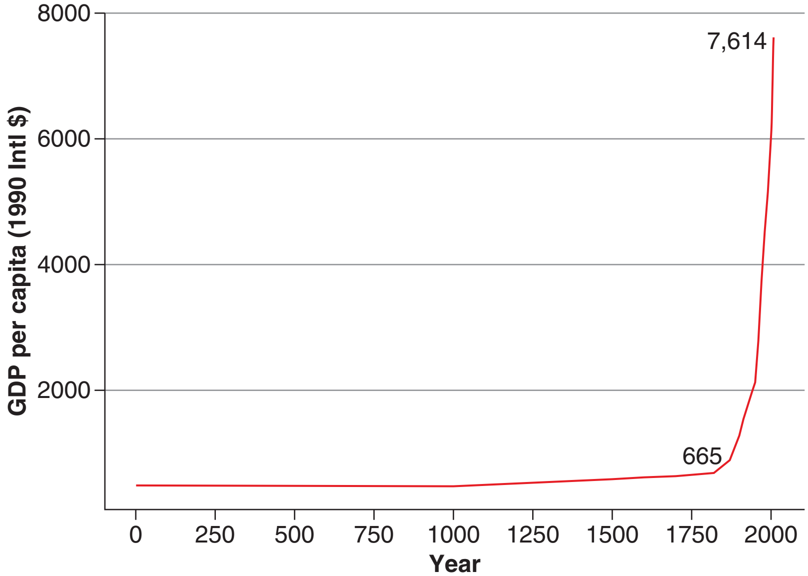

Figure 2.2 Log GDP per capita of selected countries (1820–2018)

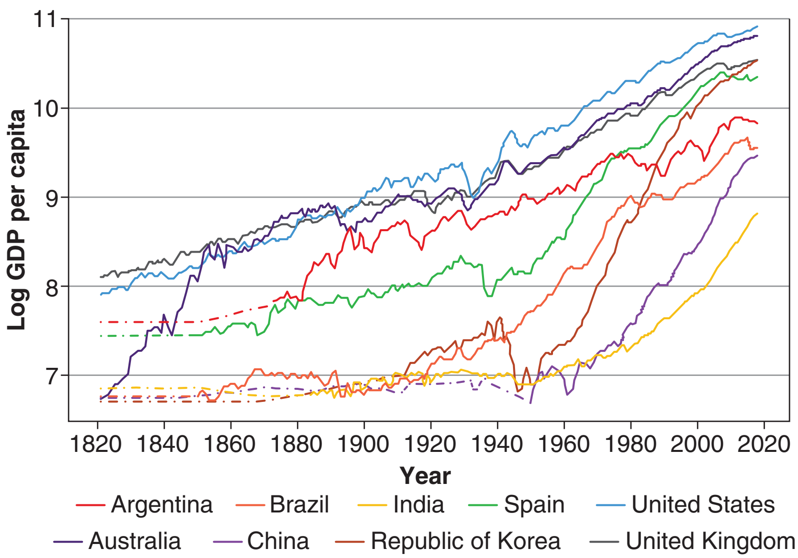

But the fundamental point of Lucas's quote is to realise that the mind-boggling differences in income per capita across countries are to a large extent due to differences in growth rates over time; and the power of exponential growth means that even relatively small differences in the latter will build into huge differences in the former. Figures 2.2 and 2.3 make this point. The richest countries

Figure 2.3 Log GDP per capita of selected countries (1960–2018)

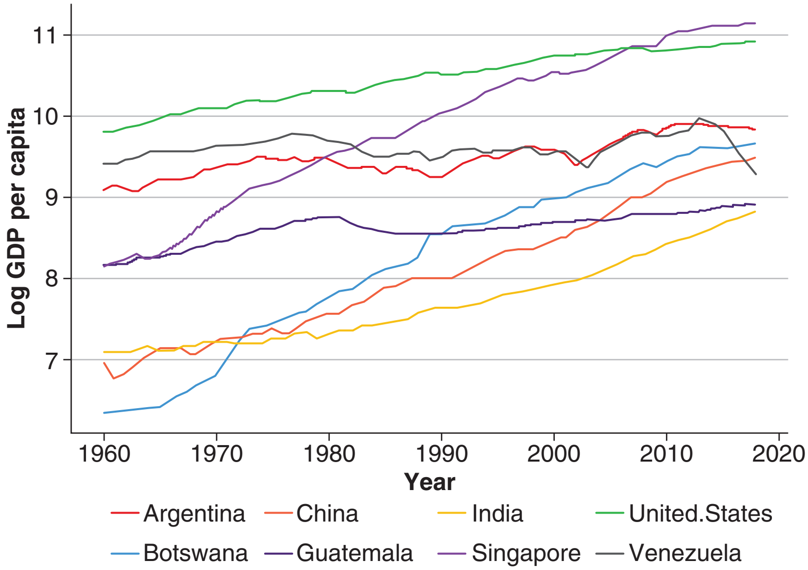

have been growing steadily over the last two centuries, and some countries have managed to converge to their income levels. Some of the performances are really stellar. Figure 2.2 shows how South Korea, with an income level that was 16% of that of the U.S. in 1940, managed to catch up in just a few decades. Today it's income is 68.5% compared to the U.S. Likewise, Spain's income in 1950 was 23% that of the U.S. Today it is 57%. At the same time other countries lagged. Argentina for example dropped from an income level that was 57% of U.S. income at the turn of the century to 33.5% today.

Figure 2.3 shows some diversity during recent times. The spectacular performances of Botswana, Singapore or, more recently, of China and India, contrast with the stagnation of Guatemala, Argentina or Venezuela. In 1960 the income of the average Motswana (as someone from Botswana is called) was only 6% as rich as the average Venezuelan. In 2018 he or she was 48% richer!

These are crucial reasons why we will spend about the initial half of this book in understanding growth. But those are not the only reasons! You may be aware that macroeconomists disagree on a lot of things; however, the issue of economic growth is one where there is much more of a consensus. It is thus helpful to start off on this relatively more solid footing. Even more importantly, the study of economic growth brings to the forefront two key ingredients of essentially all of macroeconomic analysis: general equilibrium and dynamics. First, understanding the behaviour of an entire economy requires thinking about how different markets interact and affect one another, which inevitably requires a general equilibrium approach. Second, to think seriously about how an economy evolves over time we must consider how today's choices affect tomorrow's – in other words, we must think dynamically! As such, economic growth is the perfect background upon which to develop the main methodological tools in macroeconomics: the model of intertemporal optimisation, known as the neoclassical growth model (NGM for short, also known as the Ramsey model), and the overlapping generations model (we'll call it the OLG model). A lot of what we will do later, as we explore different macroeconomic policy issues, will involve applications of these dynamic general-equilibrium tools that we will learn in the context of studying economic growth.

So, without further delay, to this we turn.

2.2 | The Kaldor facts

What are the key stylised facts about growth that our models should try to match? That there is growth in output and capital per worker with relatively stable income shares.

The modern study of economic growth starts in the post-war period and was mostly motivated by the experience of the developed world. In his classical article (Kaldor 1957), Nicolas Kaldor stated some basic facts that he observed economic growth seemed to satisfy, at least in those countries. These came to be known as the Kaldor facts, and the main challenge of growth theory as initially constituted was to account simultaneously for all these facts. But, what were these Kaldor facts? Here they are:[2]

1. *Output per worker shows continuous growth,* with no tendency to fall.
2. *The capital/output ratio is nearly constant.* (But what is capital?)
3. *Capital per worker shows continuous growth* (... follows from the other two).
4. *The rate of return on capital is nearly constant* (real interest rates are flat).
5. *Labour and capital receive constant shares of total income.*
6. *The growth rate of output per worker differs substantially across countries* (and over time, we can add, miracles and disasters).

Most of these facts have aged well. But not all of them. For example, we now know the constancy of the interest rate is not so when seen from a big historical sweep. In fact, interest rates have been on a secular downward trend that can be dated back to the 1300's (Schmelzing 2019). (Of course rates are way down now, so the question is how much lower can they go?) We will show you the data in a few pages.

In addition, in recent years, particularly since the early 1980s, the labour share has fallen significantly in most countries and industries. There is much argument in the literature as to the reasons why (see Karabarbounis and Neiman (2014) for a discussion on this) and the whole debate about income distribution trends recently spearheaded by Piketty (2014) has to do with this issue. We will come back to it shortly.

As it turns out Robert Solow established a simple model (Solow 1956) that became the first working model of economic growth.[3] Solow's contribution became the foundation of the NGM, and the backbone of modern growth theory, as it seemed to fit the Kaldor facts. Any study of growth must start with this model, reviewing what it explains – and, just as crucially, what it fails to explain.[4]

2.3 | The Solow model

We outline and solve the basic Solow model, introducing the key concepts of the neoclassical production function, the balanced growth path, transitional dynamics, dynamic inefficiency, and convergence.

Consider an economy with only two inputs: physical capital, K, and labour, L. The production function is

$$Y = F(K, L, t), \tag{2.1}$$

where Y is the flow of output produced. Assume output is a homogeneous good that can be consumed, C, or invested, I, to create new units of physical capital.

Let s be the fraction of output that is saved – that is, the *saving rate* – so that $1 - s$ is the fraction of output that is consumed. Note that $0 \leq s \leq 1$.

Assume that capital depreciates at the constant rate $\delta > 0$. The net increase in the stock of physical capital at a point in time equals gross investment less depreciation:

$$\dot{K} = I - \delta K = s \cdot F(K, L, t) - \delta K, \tag{2.2}$$

where a dot over a variable, such as $\dot{K}$, denotes differentiation with respect to time. Equation (2.2) determines the dynamics of K for a given technology and labour force.

Assume the population equals the labour force, L. It grows at a constant, exogenous rate, $\dot{L}/L = n \geq 0$.[5] If we normalise the number of people at time 0 to 1, then

$$L_t = e^{nt}. \tag{2.3}$$

where L_t is labour at time t.

If L_t is given from (2.3) and technological progress is absent, then (2.2) determines the time paths of capital, K, and output, Y. Such behaviour depends crucially on the properties of the production function, $F(\cdot)$. Apparently minor differences in assumptions about $F(\cdot)$ can generate radically different theories of economic growth.

2.3.1 | The (neoclassical) production function

For now, neglect technological progress. That is, assume that $F(\cdot)$ is independent of t. This assumption will be relaxed later. Then, the production function (2.1) takes the form

$$Y = F(K, L). \tag{2.4}$$

Assume also the following three properties are satisfied. First, for all $K > 0$ and $L > 0$, $F(\cdot)$ exhibits positive and diminishing marginal products with respect to each input:

$$\frac{\partial F}{\partial K} > 0, \qquad \frac{\partial^2 F}{\partial K^2} < 0$$
$$\frac{\partial F}{\partial L} > 0, \qquad \frac{\partial^2 F}{\partial L^2} < 0.$$

Second, $F(\cdot)$ exhibits constant returns to scale:

$$F(\lambda K, \lambda L) = \lambda \cdot F(K, L) \text{ for all } \lambda > 0.$$

Third, the marginal product of capital (or labour) approaches infinity as capital (or labour) goes to 0 and approaches 0 as capital (or labour) goes to infinity:

$$\lim_{K \to 0} \frac{\partial F}{\partial K} = \lim_{L \to 0} \frac{\partial F}{\partial L} = \infty,$$
$$\lim_{K \to \infty} \frac{\partial F}{\partial K} = \lim_{L \to \infty} \frac{\partial F}{\partial L} = 0.$$

These last properties are called *Inada conditions*.

We will refer to production functions satisfying those three sets of conditions as *neoclassical production functions*.

The condition of constant returns to scale has the convenient property that output can be written as

$$Y = F(K, L) = L \cdot F(K/L, 1) = L \cdot f(k), \tag{2.5}$$

where $k \equiv K/L$ is the capital-labour ratio, and the function $f(k)$ is defined to equal $F(k, 1)$. The production function can be written as

$$y = f(k), \tag{2.6}$$

where $y \equiv Y/L$ is per capita output.

One simple production function that satisfies all of the above and is often thought to provide a reasonable description of actual economies is the Cobb-Douglas function,

$$Y = AK^{\alpha}L^{1-\alpha}, \tag{2.7}$$

where $A > 0$ is the level of the technology, and α is a constant with $0 < \alpha < 1$. The Cobb-Douglas function can be written as

$$y = Ak^{\alpha}. \tag{2.8}$$

Note that $f'(k) = A\alpha k^{\alpha-1} > 0, f''(k) = -A\alpha(1-\alpha)k^{\alpha-2} < 0, \lim_{k\to\infty} f'(k) = 0$, and $\lim_{k\to 0} f'(k) = \infty$. In short, the Cobb-Douglas specification satisfies the properties of a neoclassical production function.

2.3.2 | The law of motion of capital

The change in the capital stock over time is given by (2.2). If we divide both sides of this equation by L, then we get

$$\dot{K}/L = s \cdot f(k) - \delta k. \tag{2.9}$$

The right-hand side contains per capita variables only, but the left-hand side does not. We can write $\dot{K}/L$ as a function of k by using the fact that

$$\dot{k} \equiv \frac{d(K/L)}{dt} = \dot{K}/L - nk, \tag{2.10}$$

where $n = \dot{L}/L$. If we substitute (2.10) into the expression for $\dot{K}/L$ then we can rearrange terms to get

$$\dot{k} = s \cdot f(k) - (n + \delta) \cdot k. \tag{2.11}$$

The term $n + \delta$ on the right-hand side of (2.11) can be thought of as the effective depreciation rate for the capital/labour ratio, $k \equiv K/L$. If the saving rate, s, were 0, then k would decline partly due to depreciation of K at the rate δ and partly due to the growth of L at the rate n.

Figure 2.4 shows the workings of (2.11). The upper curve is the production function, $f(k)$. The term $s \cdot f(k)$ looks like the production function except for the multiplication by the positive fraction s. The $s \cdot f(k)$ curve starts from the origin (because $f(0) = 0$), has a positive slope (because $f'(k) > 0$), and gets flatter as k rises (because $f''(k) < 0$). The Inada conditions imply that the $s \cdot f(k)$ curve is vertical at $k = 0$ and becomes perfectly flat as k approaches infinity. The other term in (2.11), $(n+\delta)\cdot k$, appears in Figure 2.1 as a straight line from the origin with the positive slope $n + \delta$.

Figure 2.4 Dynamics in the Solow model

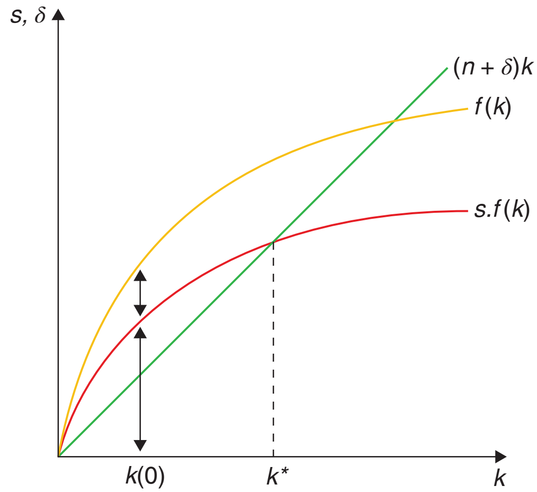

2.3.3 | Finding a balanced growth path

A *balanced growth path* (BGP) is a situation in which the various quantities grow at constant rates.[6] In the Solow model, the BGP corresponds to $\dot{k} = 0$ in (2.11).[7] We find it at the intersection of the $s \cdot f(k)$ curve with the $(n + \delta) \cdot k$ line in Figure 2.4. The corresponding value of k is denoted k^*. Algebraically, k^* satisfies the condition:

$$s \cdot f(k^*) = (n + \delta) \cdot k^*. \tag{2.12}$$

Since k is constant in the BGP, y and c are also constant at the values $y^* = f(k^*)$ and $c^* = (1 - s) \cdot f(k^*)$, respectively. Hence, in the Solow model, the per capita quantities k, y, and c do not grow in the BGP: it is a growth model without (long-term) growth!

Now, that's not quite right: the constancy of the per capita magnitudes means that the levels of variables – K, Y, and C – grow in the BGP at the rate of population growth, n. In addition, changes in the level of technology, represented by shifts of the production function, $f(\cdot)$; in the saving rate, s; in the rate of population growth, n; and in the depreciation rate, δ; all have effects on the per capita *levels* of the various quantities in the BGP.

We can illustrate the results for the case of a Cobb-Douglas production function. The capital/labour ratio on the BGP is determined from (2.12) as

$$k^* = \left(\frac{sA}{n + \delta}\right)^{\frac{1}{1-\alpha}}. \tag{2.13}$$

Note that, as we saw graphically for a more general production function $f(k)$, k^* rises with the saving rate, s, and the level of technology, A, and falls with the rate of population growth, n, and the depreciation rate, δ. Output per capita on the BGP is given by

$$y^* = A^{\frac{1}{1-\alpha}} \cdot \left(\frac{s}{n + \delta}\right)^{\frac{\alpha}{1-\alpha}}. \tag{2.14}$$

Thus, y^* is a positive function of s and A and a negative function of n and δ.

2.3.4 | Transitional dynamics

Moreover, the Solow model does generate growth in the transition to the BGP. To see the implications in this regard, note that dividing both sides of (2.11) by k implies that the growth rate of k is given by

$$\gamma_k \equiv \frac{\dot{k}}{k} = \frac{s \cdot f(k)}{k} - (n + \delta). \tag{2.15}$$

Equation (2.15) says that γ_k equals the difference between two terms, $s \cdot f(k)/k$ and $(n + \delta)$ which we plot against k in Figure 2.5. The first term is a downward-sloping curve, which asymptotes to infinity at $k = 0$ and approaches 0 as k tends to infinity. The second term is a horizontal line crossing the vertical axis at $n + \delta$. The vertical distance between the curve and the line equals the growth rate of capital per person, and the crossing point corresponds to the BGP. Since $n + \delta > 0$ and $s \cdot f(k)/k$ falls monotonically from infinity to 0, the curve and the line intersect once and only once. Hence (except for the trivial solution $k^* = 0$, where capital stays at zero forever), the BGP capital-labour ratio $k^* > 0$ exists and is unique.

Note also that output moves according to

$$\frac{\dot{y}}{y} = \alpha \frac{\dot{k}}{k} = \alpha \gamma_k. \tag{2.16}$$

A formal treatment of dynamics follows. From (2.11) one can calculate

$$\frac{d\dot{k}}{dk} = s \cdot f'(k) - (n + \delta). \tag{2.17}$$

We want to study dynamics in the neighbourhood of the BGP, so we evaluate this at k^*:

$$\left.\frac{d\dot{k}}{dk}\right|_{k=k^*} = s \cdot f'(k^*) - (n + \delta). \tag{2.18}$$

Figure 2.5 Dynamics in the Solow model again

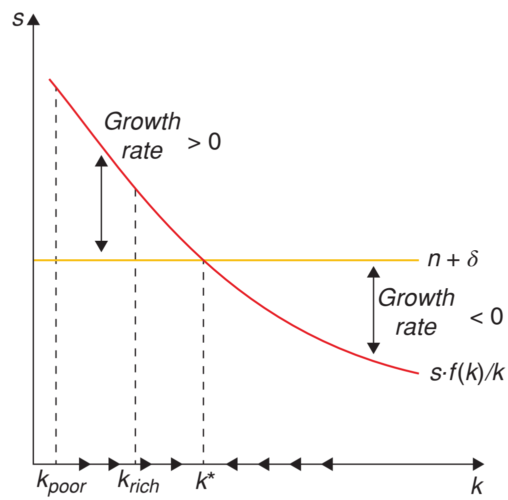

The capital stock will converge to its BGP if $\dot{k} > 0$ when $k < k^*$ and $\dot{k} < 0$ when $k > k^*$. Hence, this requires that $\left.\frac{d\dot{k}}{dk}\right|_{k=k^*} < 0$.

In the Cobb-Douglas case the condition is simple. Note that

$$\left.\frac{d\dot{k}}{dk}\right|_{k=k^*} = s \cdot A\alpha \left(\frac{sA}{n+\delta}\right)^{-1} - (n+\delta) = (n+\delta)(\alpha - 1) \tag{2.19}$$

so that $\left.\frac{d\dot{k}}{dk}\right|_{k=k^*} < 0$ requires $\alpha < 1$. That is, reaching the BGP requires diminishing returns.

With diminishing returns, when k is relatively low, the marginal product of capital, $f'(k)$, is relatively high. By assumption, households save and invest a constant fraction, s, of this product. Hence, when k is relatively low, the marginal return to investment, $s \cdot f'(k)$, is relatively high. Capital per worker, k, effectively depreciates at the constant rate $n + \delta$. Consequently, the growth of capital, $\dot{k}$, is also relatively high. In fact, for $k < k^*$ it is positive. Conversely, for $k > k^*$ it is negative.

2.3.5 | Policy experiments

Suppose that the economy is initially on a BGP with capital per person k_1^*. Imagine that the government then introduces some policy that raises the saving rate permanently from s_1 to a higher value s_2. Figure 2.6 shows that the $s \cdot f(k)/k$ schedule shifts to the right. Hence, the intersection with the $n + \delta$ line also shifts to the right, and the new BGP capital stock, k_2^*, exceeds k_1^*. An increase in the saving rate generates temporarily positive per capita growth rates. In the long run, the levels of k and y are permanently higher, but the per capita growth rates return to 0.

A permanent improvement in the level of the technology has similar, temporary effects on the per capita growth rates. If the production function, $f(k)$, shifts upward in a proportional manner, then the

Figure 2.6 The effects of an increase in the savings rate

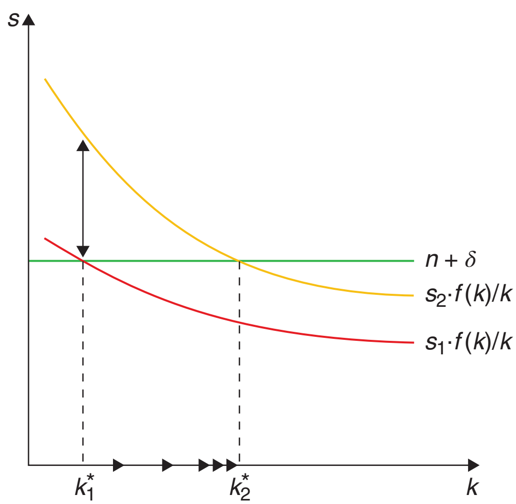

$s \cdot f(k)/k$ curve shifts upward, just as in Figure 2.6. Hence, γ_k again becomes positive temporarily. In the long run, the permanent improvement in technology generates higher levels of k and y, but no changes in the per capita growth rates.

2.3.6 | Dynamic inefficiency

For a given production function and given values of n and δ, there is a unique BGP value $k^* > 0$ for each value of the saving rate, s. Denote this relation by $k^*(s)$, with $dk^*(s)/ds > 0$. The level of per capita consumption on the BGP is $c^* = (1-s) \cdot f\left[k^*(s)\right]$. We know from (2.12) that $s \cdot f(k^*) = (n+\delta) \cdot k^*$; hence we can write an expression for c^*as

$$c^*(s) = f\left[k^*(s)\right] - (n+\delta) \cdot k^*. \tag{2.20}$$

Figure 2.7 shows the relation between $c^*$and s that is implied by (2.20). The quantity $c^*$is increasing in s for low levels of s and decreasing in s for high values of s. The quantity c^* attains its maximum when the derivative vanishes, that is, when $\left[f'(k^*) - (n+\delta)\right] \cdot dk^*/ds = 0$. Since $dk^*/ds > 0$, the term in brackets must equal 0. If we denote the value of k^* by k_g that corresponds to the maximum of c^*, then the condition that determines k_g is

$$f'\left(k_g\right) = (n+\delta). \tag{2.21}$$

The corresponding savings rate can be denoted as s_g, and the associated level of per capita consumption on the BGP is given by $c_g = f\left(k_g\right) - (n+\delta) \cdot k_g$ and is is called the "golden rule" consumption rate.

If the savings rate is greater than that, then it is possible to increase consumption on the BGP, and also over the transition path. We refer to such a situation, where everyone could be made better off by an alternative allocation, as one of *dynamic inefficiency*. In this case, this dynamic inefficiency is brought about by oversaving: everyone could be made better off by choosing to save less and consume more. But this naturally begs the question: why would anyone pass up this opportunity? Shouldn't we

Figure 2.7 Feasible consumption

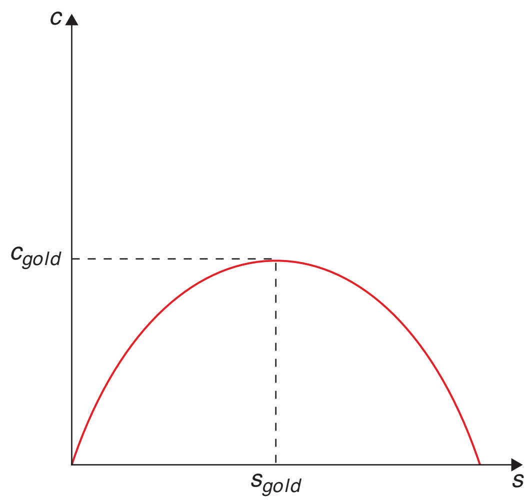

think of a better model of how people make their savings decisions? We will see about that in the next chapter.

2.3.7 | Absolute and conditional convergence

Equation (2.15) implies that the derivative of γ_k with respect to k is negative:

$$\partial\gamma_k/\partial k = \frac{s}{k}\left[f'(k) - \frac{f(k)}{k}\right] < 0. \tag{2.22}$$

Other things equal, smaller values of k are associated with larger values of γ_k. Does this result mean that economies with lower capital per person tend to grow faster in per capita terms? Is there *convergence* across economies?

We have seen above that economies that are structurally similar in the sense that they have the same values of the parameters s, n, and δ and also have the same production function, $F(\cdot)$, have the same BGP values k^* and y^*. Imagine that the only difference among the economies is the initial quantity of capital per person, $k(0)$. The model then implies that the less-advanced economies – with lower values of $k(0)$ and $y(0)$ – have higher growth rates of k. This hypothesis is known as *conditional convergence*: within a group of structurally similar economies (i.e. with similar values for s, n, and δ and production function, $F(\cdot)$), poorer economies will grow faster and catch up with the richer one. This hypothesis does seem to match the data – think about how poorer European countries have grown faster, or how the U.S. South has caught up with the North, over the second half of the 20th century.

An alternative, stronger hypothesis would posit simply that poorer countries would grow faster without conditioning on any other characteristics of the economies. This is referred to as *absolute convergence*, and does not seem to fit the data well.[8] Then again, the Solow model does *not* predict absolute convergence!

2.4 | Can the model account for income differentials?

We have seen that the Solow model does not have growth in per capita income in the long run. But can it help us understand income differentials?

We will tackle the empirical evidence on economic growth at a much greater level of detail later on. However, right now we can ask whether the simple Solow model can account for the differences in income levels that are observed in the world. According to the World Bank's calculations, the range of 2020 PPP income levels vary from \$ 138,000 per capita in Qatar or \$80,000 in Norway, all the way down to \$ 700 in Burundi. Can the basic Solow model explain this difference in income per capita of a factor of more than 100 times or even close to 200 times?

In order to tackle this question we start by remembering what output is supposed to be on the BGP:

$$y^* = A^{\frac{1}{1-\alpha}}\left(\frac{s}{n+\delta}\right)^{\frac{\alpha}{1-\alpha}}. \tag{2.23}$$

Assuming $A = 1$ and $n = 0$ this simplifies to:

$$y^* = \left(\frac{s}{\delta}\right)^{\frac{\alpha}{1-\alpha}}. \tag{2.24}$$

The ability of the Solow model to explain these large differences in income (in the BGP), as can be seen from the expressions above, will depend critically on the value of α.

$$\text{If} \begin{cases} \alpha = \frac{1}{3} \text{ then } \frac{\alpha}{1-\alpha} = \frac{1/3}{2/3} = \frac{1}{2} \\ \alpha = \frac{1}{2} \text{ then } \frac{\alpha}{1-\alpha} = \frac{1/2}{1/2} = 1 \\ \alpha = \frac{2}{3} \text{ then } \frac{\alpha}{1-\alpha} = \frac{2/3}{1/3} = 2. \end{cases}$$

The standard (rough) estimate for the capital share is $\frac{1}{3}$. Parente and Prescott (2002), however, claim that the capital share in GDP is much larger than usually accounted for because there are large intangible capital assets. In fact, they argue that the share of investment in GDP is closer to two-thirds rather than the more traditional one-third. The reasons for the unaccounted investment are (their estimates of the relevance of each in parenthesis):

1. Repair and maintenance (5% of GDP)
2. R&D (3% of GDP) multiplied by three (i.e. 9% of GDP) to take into account perfecting the manufacturing process and launching new products (the times three is not well substantiated)
3. Investment in software (3% of GDP)
4. Firms investment in organisation capital. (They think 12% is a good number.)
5. Learning on the job and training (10% of GDP)
6. Schooling (5% of GDP)

They claim all this capital has a return and that it accounts for about 56% of total GDP!

At any rate, using the equation above:

$$\frac{y_1}{y_2} = \frac{\left(\frac{s_1}{\delta}\right)^{\frac{\alpha}{1-\alpha}}}{\left(\frac{s_2}{\delta}\right)^{\frac{\alpha}{1-\alpha}}} = \left(\frac{s_1}{s_2}\right)^{\frac{\alpha}{1-\alpha}}, \tag{2.25}$$

which we can use to estimate income level differences.

	$\left(\frac{y_1}{y_2} - 1\right) * 100$		
$\frac{s_1}{s_2}$	$\alpha = \frac{1}{3}$	$\alpha = \frac{1}{2}$	$\alpha = \frac{2}{3}$
1	0%	0%	0%
1.5	22%	50%	125%
2	41%	100%	300%
3	73%	200%	800%

But even the 800% we get using the two-thirds estimate seems to be way too low relative to what we see in the data.

Alternatively, the differences in income may come from differences in total factor productivity (TFP), as captured by A. The question is: how large do these differences need to be to explain the output differentials? Recall from (2.23) that

$$y^* = A^{\frac{1}{1-\alpha}} \left(\frac{s}{n+\delta}\right)^{\frac{\alpha}{1-\alpha}}. \tag{2.26}$$

So if $\alpha = 2/3$, as suggested by Parente and Prescott (2002), then $A^{\frac{1}{1-\alpha}} = A^{\frac{1}{1/3}} = A^3$. Now, let's forget about s, δ, n (for example, by assuming they are the same for all countries), and just focus on differences in A. Notice that if TFP is $1/3$, of the level in the other country, this indicates that the income level is then $1/27$.

Parente and Prescott (2002) use this to estimate, for a group of countries, how much productivity would have to differ (relative to the United States) for us to replicate observed relative incomes over the period 1950–1988:

Country	Relative Income		Relative TFP
UK	60%	→	86%
Colombia	22%	→	64%
Paraguay	16%	→	59%
Pakistan	10%	→	51%

These numbers appear quite plausible, so the message is that the Solow model requires substantial cross-country differences in productivity to approximate existing cross-country differences in income. This begs the question of what makes productivity so different across countries, but we will come back to this later.

2.5 | The Solow model with exogenous technological change

We have seen that the Solow model does not have growth in per capita income in the long run. But that changes if we allow for technological change.

Allow now the productivity of factors to change over time. In the Cobb-Douglas case, this means that A increases over time. For simplicity, suppose that $\dot{A}/A = a > 0$. Out of the BGP, output then evolves according to

$$\frac{\dot{y}}{y} = \frac{\dot{A}}{A} + \alpha\frac{\dot{k}}{k} = a + \alpha\gamma_k. \tag{2.27}$$

On the BGP, where k is constant,

$$\frac{\dot{y}}{y} = a. \tag{2.28}$$

This is a strong prediction of the Solow model: in the long run, technological change is the only source of growth in per capita income.

Let's now embed this improvement in technology or efficiency in workers. We can define labour input as broader than just bodies, we could call it now human capital defined by

$$E_t = L_t \cdot e^{\lambda t} = L_0 \cdot e^{(\lambda+n)t}, \tag{2.29}$$

where E is the amount of labor in efficiency units. The production function is

$$Y = F\left(K_t, E_t\right). \tag{2.30}$$

To put it in per capita efficiency terms, we define

$$k = \frac{K}{E}. \tag{2.31}$$

So

$$\frac{\dot{k}}{k} = \frac{\dot{K}}{K} - \frac{\dot{E}}{E} = \frac{sy}{k} - \delta - n - \lambda, \tag{2.32}$$

$$\frac{\dot{k}}{k} = \frac{sf(k)}{k} - \delta - n - \lambda, \tag{2.33}$$

$$\dot{k} = sf(k) - (\delta + n + \lambda)\, k. \tag{2.34}$$

For $\dot{k} = 0$

$$\frac{sf(k)}{k} = (\delta + n + \lambda)\,. \tag{2.35}$$

On the BGP $\dot{k} = 0$, so

$$\frac{\dot{K}}{K} = \frac{\dot{E}}{E} = n + \lambda = \frac{\dot{Y}}{Y}. \tag{2.36}$$

But then

$$\frac{\dot{\left(\frac{Y}{L}\right)}}{\frac{Y}{L}} = \frac{\dot{Y}}{Y} - \frac{\dot{L}}{L} = \lambda \tag{2.37}$$

Notice that in this equilibrium income per person grows even on the BGP, and this accounts for all six Kaldor facts.

2.6 | What have we learned?

The Solow model shows that capital accumulation by itself cannot sustain growth in per capita income in the long run. This is because accumulation runs into diminishing marginal returns. At some point the capital stock becomes large enough – and its marginal product correspondingly small enough – that a given savings rate can only provide just enough new capital to replenish ongoing depreciation and increases in labour force. Alternatively, if we introduce exogenous technological change that increases productivity, we can generate long-run growth in income per capita, but we do not really explain it. In fact, any differences in long-term growth rates come from exogenous differences in the rate of technological change – we are not explaining those differences, we are just assuming them! As a result, nothing within the model tells you what policy can do about growth in the long run.

That said, we do learn a lot about growth in the transition to the long run, about differences in income levels, and how policy can affect those things. There are clear lessons about: (i) convergence – the model predicts conditional convergence; (ii) dynamic inefficiency – it is possible to save too much in this model; and (iii) long-run differences in income – they seem to have a lot to do with differences in productivity.

Very importantly, the model also points at the directions we can take to try and understand long-term growth. We can have a better model of savings behaviour: how do we know that individuals will save what the model says they will save? And, how does that relate to the issue of dynamic inefficiency?

We can look at different assumptions about technology: maybe we can escape the shackles of diminishing returns to accumulation? Or can we think more carefully about how technological progress comes about?

These are the issues that we will address over the next few chapters.

Notes

[1] Lucas's words hold up very well more than three decades later, in spite of some evidently dated examples.

[2] Once we are done with our study of economic growth, you can check the "new Kaldor facts" proposed by Jones and Romer (2010), which update the basic empirical regularities based on the progress over the subsequent half-century or so.

[3] For those of you who are into the history of economic thought, at the time the framework to study growth was the so-called Harrod-Domar model, due to the independent contributions of (you probably guessed it...) Harrod (1939) and Domar (1946). It assumed a production function with perfect complementarity between labour and capital ("Leontieff", as it is known to economists), and predicted that an economy would generate increasing unemployment of either labour or capital, depending on whether it saved a little or a lot. As it turns out, that was not a good description of the real world in the post-war period.

[4] Solow eventually got a Nobel prize for his trouble, in 1987 – also for his other contributions to the study of economic growth, to which we will return. An Australian economist, Trevor Swan, also published an independently developed paper with very similar ideas at about the same time, which is why sometimes the model is referred to as the Solow-Swan model. He did not get a Nobel prize.

[5] We will endogenise population growth in Chapter 10, when discussing unified growth theory.

[6] The BGP is often referred to as a "steady state", borrowing terminology from classical physics. We have noticed that talk of "steady state" tends to lead students to think of a situation where variables are not growing at all. The actual definition refers to constant growth rates, and it is only in certain cases and for certain variables, as we will see, that this constant rate happens to be zero.

[7] You should try to show mathematically from (2.11) that, with a neoclassical production function, the only way we can have a constant growth rate $\frac{\dot{k}}{k}$ is to have $\dot{k} = 0$.

[8] Or does it? More recently, Kremer et al. (2021) have argued that there has been a move towards absolute convergence in the data in the 21st century... Stay tuned!

References

Domar, E. D. (1946). Capital expansion, rate of growth, and employment. *Econometrica*, 137–147.

Harrod, R. F. (1939). An essay in dynamic theory. *The Economic Journal*, *49*(193), 14–33.

Jones, C. I. & Romer, P. M. (2010). The new Kaldor facts: Ideas, institutions, population, and human capital. *American Economic Journal: Macroeconomics*, *2*(1), 224–245.

Kaldor, N. (1957). A model of economic growth. *Economic Journal*, *67*(268), 591–624.

Karabarbounis, L. & Neiman, B. (2014). The global decline of the labor share. *The Quarterly Journal of Economics*, *129*(1), 61–103.

Kremer, M., Willis, J., & You, Y. (2021). Converging to convergence. *NBER Macro Annual 2021*. https://www.nber.org/system/files/chapters/c14560/c14560.pdf.

Lucas Jr., R. (1988). On the mechanics of economic development. *Journal of Monetary Economics*, *22*, 3–42.

Parente, S. L. & Prescott, E. C. (2002). *Barriers to riches*. MIT Press.

Piketty, T. (2014). *Capital in the twenty-first century*. Harvard University Press.

Schmelzing, P. (2019). Eight centuries of global real rates, r-g, and the 'suprasecular' decline, 1311–2018, Available at SSRN: https://ssrn.com/abstract=3485734 or http://dx.doi.org/10.2139/ssrn.3485734.

Solow, R. M. (1956). A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, *70*(1), 65–94.

CHAPTER 3

The neoclassical growth model

3.1 | The Ramsey problem

We will solve the optimal savings problem underpinning the Neoclassical Growth Model, and in the process introduce the tools of dynamic optimisation we will use throughout the book. We will also encounter, for the first time, the most important equation in macroeconomics: the Euler equation.

$$\frac{\dot{c}_t}{c_t} = \sigma\left[f'\left(k_t\right) - \rho\right]$$

We have seen the lessons and shortcomings of the basic Solow model. One of its main assumptions, as you recall, was that the savings rate was constant. In fact, there was no optimisation involved in that model, and welfare statements are hard to make in that context. This is, however, a very rudimentary assumption for an able policy maker who is in possession of the tools of dynamic optimisation. Thus we tackle here the challenge of setting up an optimal program where savings is chosen to maximise intertemporal welfare.

As it turns out, British philosopher and mathematician Frank Ramsey, in one of the two seminal contributions he provided to economics before dying at the age of 26, solved this problem in 1928 (Ramsey (1928)).[1] The trouble is, he was so ahead of his time that economists would only catch up in the 1960s, when David Cass and Tjalling Koopmans independently revived Ramsey's contribution.[2] (That is why this model is often referred to either as the Ramsey model or the Ramsey-Cass-Koopmans model.) It has since become ubiquitous and, under the grand moniker of Neoclassical Growth Model (NGM), it is the foremost example of the type of *dynamic general equilibrium* model upon which the entire edifice of modern macroeconomics is built.

To make the problem manageable, we will assume that there is one representative household, all of whose members are both consumer and producer, living in a closed economy (we will lift this assumption in the next chapter). There is one good and no government. Each consumer in the representative household lives forever, and population growth is $n > 0$ as before. All quantities in small-case letters are per capita. Finally, we will look at the problem as solved by a benevolent central planner who maximises the welfare of that representative household, and evaluates the utility of future consumption at a discounted rate.

At this point, it is worth stopping and thinking about the model's assumptions. By now you are already used to outrageously unrealistic assumptions, but this may be a little too much. People

How to cite this book chapter:
Campante, F., Sturzenegger, F. and Velasco, A. 2021. *Advanced Macroeconomics: An Easy Guide.* Ch. 3. 'The neoclassical growth model', pp. 23–40. London: LSE Press.
DOI: https://doi.org/10.31389/lsepress.ame.c

obviously do not live forever, they are not identical, and what's this business of a benevolent central planner? Who are they? Why would they discount future consumption? Let us see why we use these shortcuts:

1. We will look at the central planner's problem, as opposed to the decentralised equilibrium, because it is easier and gets us directly to an efficient allocation. We will show that, under certain conditions, it provides the same results as the decentralised equilibrium. This is due to the so-called welfare theorems, which you have seen when studying microeconomics, but which we should perhaps briefly restate here:
   a. A competitive equilibrium is Pareto Optimal.
   b. All Pareto Optimal allocations can be decentralised as a competitive equilibrium under some convexity assumptions. Convexity of production sets means that we cannot have increasing returns to scale. (If we do, we need to depart from competitive markets.)
2. There's only one household? Certainly this is not very realistic, but it is okay if we think that typically people react similarly (not necessarily identically) to the parameters of the model. Specifically, do people respond similarly to an increase in the interest rate? If you think they do, then the assumption is okay.
3. Do all the people have the same utility function? Are they equal in all senses? Again, as above, not really. But, we believe they roughly respond similarly to basic tradeoffs. In addition, as shown by Caselli and Ventura (2000), one can incorporate a lot of sources of heterogeneity (namely, individuals can have different tastes, skills, initial wealth) and still end up with a representative household representation, as long as that heterogeneity has a certain structure. The assumption also means that we are, for the most part, ignoring distributional concerns, but that paper also shows that a wide range of distributional dynamics are compatible with that representation. (We will also raise some points about inequality as we go along.)
4. Do they live infinitely? Certainly not, but it does look like we have some intergenerational links. Barro (1974) suggests an individual who cares about the utility of their child: $u\left(c_t\right) + \beta V\left[u\left(c_{child}\right)\right]$. If that is the case, substituting recursively gives an intertemporal utility of the sort we have posited. And people do think about the future.
5. Why do we discount future utility? To some extent it is a revealed preference argument: interest rates are positive and this only makes sense if people value more today's consumption than tomorrow's, which is what we refer to when we speak of discounting the future. On this you may also want to check Caplin and Leahy (2004), who argue that a utility such as that in (3.1) imposes a sort of tyranny of the present: past utility carries no weight, whereas future utility is discounted. But does this make sense from a planner's point of view? Would this make sense from the perspective of tomorrow? In fact, Ramsey argued that it was unethical for a central planner to discount future utility.[3]

Having said that, let's go solve the problem.

3.1.1 | The consumer's problem

The utility function is[4]

$$\int_0^\infty u(c_t)e^{nt}e^{-\rho t}dt, \tag{3.1}$$

where c_t denotes consumption per capita and $\rho\,(> n)$ is the rate of time preference.[5] Assume $u'(c_t) > 0$, $u''(c_t) \leq 0$, and Inada conditions are satisfied.

3.1.2 | The resource constraint

The resource constraint of the economy is

$$\dot{K}_t = Y_t - C_t = F\left(K_t, L_t\right) - C_t, \tag{3.2}$$

with all variables as defined in the previous chapter. (Notice that for simplicity we assume there is no depreciation.) In particular, $F\left(K_t, L_t\right)$ is a neoclassical production function – hence neoclassical growth model. You can think of household production: household members own the capital and they work for themselves in producing output. Each member of the household inelastically supplies one unit of labour per unit of time.

This resource constraint is what makes the problem truly dynamic. The capital stock in the future depends on the choices that are made in the present. As such, the capital stock constitutes what we call the *state variable* in our problem: it describes the state of our dynamic system at any given point in time. The resource constraint is what we call the *equation of motion*: it characterises the evolution of the state variable over time. The other key variable, consumption, is what we call the *control variable*: it is the one variable that we can directly choose. Note that the control variable is jumpy: we can choose whatever (feasible) value for it at any given moment, so it can vary discontinuously. However, the state variable is sticky: we cannot change it discontinuously, but only in ways that are consistent with the equation of motion.

Given the assumption of constant returns to scale, we can express this constraint in per capita terms, which is more convenient. Dividing (3.2) through by L we get

$$\frac{\dot{K}_t}{L_t} = F\left(k_t, 1\right) - c_t = f\left(k_t\right) - c_t, \tag{3.3}$$

where $f(.)$ has the usual properties. Recall

$$\frac{\dot{K}_t}{L_t} = \dot{k}_t + nk_t. \tag{3.4}$$

Combining the last two equations yields

$$\dot{k}_t = f\left(k_t\right) - nk_t - c_t, \tag{3.5}$$

which we can think of as the relevant budget constraint. This is the final shape of the equation of motion of our dynamic problem, describing how the variable responsible for the dynamic nature of the problem – in this case the per capita capital stock k_t – evolves over time.

3.1.3 | Solution to consumer's problem

The household's problem is to maximise (3.1) subject to (3.5) for given k_0. If you look at our mathematical appendix, you will learn how to solve this, but it is instructive to walk through the steps here, as they have intuitive interpretations. You will need to set up the (current value) Hamiltonian for the problem, as follows:

$$H = u(c_t)e^{nt} + \lambda_t\left[f\left(k_t\right) - nk_t - c_t\right]. \tag{3.6}$$

Recall that c is the control variable (jumpy), and k is the state variable (sticky), but the Hamiltonian brings to the forefront another variable: λ, the *co-state variable*. It is the multiplier associated with the intertemporal budget constraint, analogously to the Lagrange multipliers of static optimisation.

Just like its Lagrange cousin, the co-state variable has an intuitive economic interpretation: it is the marginal value as of time t (i.e. the current value) of an additional unit of the state variable (capital, in this case). It is a (shadow) price, which is also jumpy.

First-order conditions (FOCs) are

$$\frac{\partial H}{\partial c_t} = 0 \Rightarrow u'(c_t)e^{nt} - \lambda_t = 0, \tag{3.7}$$

$$\dot{\lambda}_t = -\frac{\partial H}{\partial k_t} + \rho\lambda_t \Rightarrow \dot{\lambda}_t = -\lambda_t \left[f'\left(k_t\right) - n\right] + \rho\lambda_t, \tag{3.8}$$

$$\lim_{t\to\infty}\left(k_t\lambda_t e^{-\rho t}\right) = 0. \tag{3.9}$$

What do these optimality conditions mean? First, (3.7) should be familiar from static optimisation: differentiate with respect to the control variable, and set that equal to zero. It makes sure that, at any given point in time, the consumer is making the optimal decision – otherwise, she could obviously do better... The other two are the ones that bring the dynamic aspects of the problem to the forefront. Equation (3.9) is known as the transversality condition (TVC). It means, intuitively, that the consumer wants to set the optimal path for consumption such that, in the "end of times" (at infinity, in this case), they are left with no capital. (As long as capital has a positive value as given by λ, otherwise they don't really care...) If that weren't the case, I would be "dying" with valuable capital, which I could have used to consume a little more over my lifetime.

Equation (3.8) is the FOC with respect to the state variable, which essentially makes sure that at any given point in time the consumer is leaving the optimal amount of capital for the future. But how so? As it stands, it has been obtained mechanically. However, it is much nicer when we derive it purely from economic intuition. Note that we can rewrite it as follows:

$$\frac{\dot{\lambda}_t}{\lambda_t} = \rho - \left(f'\left(k_t\right) - n\right) \Rightarrow \rho + n = \frac{\dot{\lambda}_t}{\lambda_t} + f'\left(k_t\right). \tag{3.10}$$

This is nothing but an arbitrage equation for a typical asset price, where in this case the asset is the capital stock of the economy. Such arbitrage equations state that the opportunity cost of holding the asset (ρ in this case), equals its rate of return, which comprises the dividend yield ($f'(k_t) - n$) plus whatever capital gain you may get from holding the asset ($\frac{\dot{\lambda}_t}{\lambda_t}$). If the opportunity cost were higher (resp. lower), you would not be in an optimal position: you should hold less (resp. more) of the asset. We will come back to this intuition over and over again.

3.1.4 | The balanced growth path and the Euler equation

We are ultimately interested in the dynamic behaviour of our control and state variables, c_t and k_t. How can we turn our FOCs into a description of that behaviour (preferably one that we can represent graphically)? We start by taking (3.7) and differentiating both sides with respect to time:

$$u''(c_t)\dot{c}_t e^{nt} + nu'(c_t)e^{nt} = \dot{\lambda}_t. \tag{3.11}$$

Divide this by (3.7) and rearrange:

$$\frac{u''(c_t)c_t}{u'(c_t)}\frac{\dot{c}_t}{c_t} = \frac{\dot{\lambda}_t}{\lambda_t} - n. \tag{3.12}$$

Next, define

$$\sigma \equiv -\frac{u'(c_t)}{u''(c_t)c_t} > 0 \tag{3.13}$$

as the elasticity of intertemporal substitution in consumption.[6] Then, (3.12) becomes

$$\frac{\dot{c}_t}{c_t} = -\sigma\left(\frac{\dot{\lambda}_t}{\lambda_t} - n\right). \tag{3.14}$$

Finally, using (3.10) in (3.14) we obtain

$$\frac{\dot{c}_t}{c_t} = \sigma\left[f'\left(k_t\right) - \rho\right]. \tag{3.15}$$

This dynamic optimality condition is known as the Ramsey rule (or Keynes-Ramsey rule), and in a more general context it is referred to as the *Euler equation*. It may well be the most important equation in all of macroeconomics: it encapsulates the essence of the solution to any problem that trades off today versus tomorrow.[7]

But what does it mean intuitively? Think about it in these terms: if the consumer postpones the enjoyment of one unit of consumption to the next instant, it will be incorporated into the capital stock, and thus yield an extra $f'(\cdot)$. However, this will be worth less, by a factor of ρ. They will only consume more in the next instant (i.e. $\frac{\dot{c}_t}{c_t} > 0$) if the former compensates for the latter, as mediated by their proclivity to switch consumption over time, which is captured by the elasticity of intertemporal substitution, σ. Any dynamic problem we will see from now on involves some variation upon this general theme: the optimal growth rate trades off the rate of return of postponing consumption (i.e. investment) against the discount rate.

Mathematically speaking, equations (3.5) and (3.15) constitute a system of two differential equations in two unknowns. These plus the initial condition for capital and the TVC fully characterise the dynamics of the economy: once we have c_t and k_t, we can easily solve for any remaining variables of interest.

To make further progress, let us characterise the BGP of this economy. Setting (3.5) equal to zero yields

$$c^* = f(k^*) - nk^*, \tag{3.16}$$

which obviously is a hump-shaped function in c, k space. The dynamics of capital can be understood with reference to this function (Figure 3.1): for any given level of capital, if consumption is higher (resp. lower) than the BGP level, this means that the capital stock will decrease (resp. increase).

By contrast, setting (3.15) equal to zero yields

$$f'\left(k^*\right) = \rho. \tag{3.17}$$

This equation pins down the level of the capital stock on the BGP, and the dynamics of consumption can be seen in Figure 3.2: for any given level of consumption, if the capital stock is below (resp. above) its BGP level, then consumption is increasing (resp. decreasing). This is because the marginal product of capital will be relatively high (resp. low).

Expressions (3.16) and (3.17) together yield the values of consumption and the capital stock (both per-capita) in the BGP, as shown in Figure 3.3. This already lets us say something important about the behaviour of this economy. Let's recall the concept of the *golden rule*, from our discussion of the

Figure 3.1 Dynamics of capital

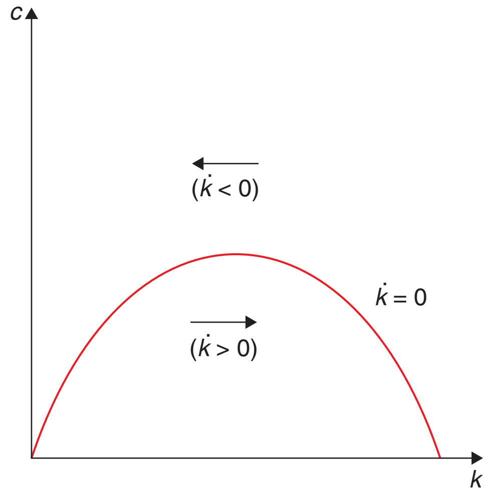

Figure 3.2 Dynamics of consumption

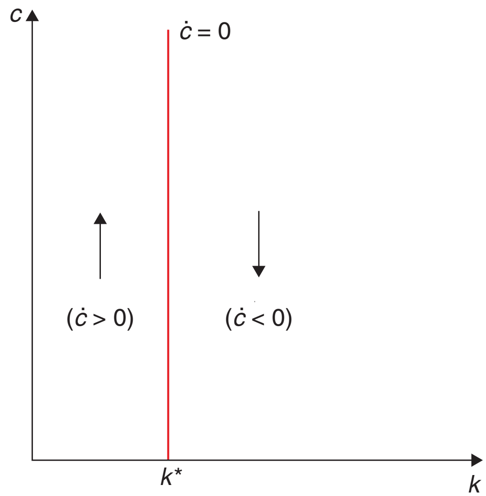

Solow model: the maximisation of per-capita consumption on the BGP. From (3.16) we see that this is tantamount to setting

$$\frac{\partial c^*}{\partial k^*} = f'\left(k_G^*\right) - n = 0 \Rightarrow f'\left(k_G^*\right) = n. \tag{3.18}$$

(Recall here we have assumed the depreciation rate is zero ($\delta = 0$).) If we compare this to (3.17), we see that the the optimal BGP level of capital per capita is *lower* than in the golden rule from the Solow model. (Recall the properties of the neoclassical production function, and that we assume $\rho > n$.)

Because of this comparison, (3.17) is sometimes known as the *modified golden rule*. Why does optimality require that consumption be lower on the BGP than what would be prescribed by the Solow

Figure 3.3 Steady state

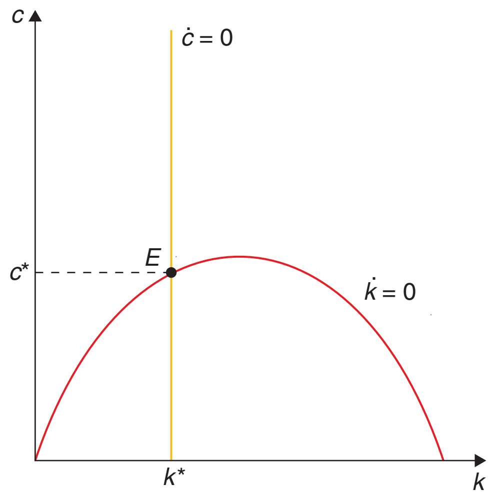

golden rule? Because future consumption is discounted, it is not optimal to save so much that BGP consumption is maximised – it is best to consume more along the transition to the BGP. Keep in mind that it is (3.17), not (3.18), that describes the optimal allocation. The kind of oversaving that is possible in the Solow model disappears once we consider optimal savings decisions.

Now, you may ask: is it the case then that this type of oversaving is not an issue in practice (or even just in theory)? Well, we will return to this issue in Chapter 8. For now, we can see how the question of dynamic efficiency relates to issues of inequality.

3.1.5 | A digression on inequality: Is Piketty right?

It turns out that we can say something about inequality in the context of the NGM, even though the representative agent framework does not address it directly. Let's start by noticing that, as in the Solow model, on the BGP output grows at the rate n of population growth (since capital and output per worker are constant). In addition, once we solve for the decentralised equilibrium, which we sketch in Section 2 below, we will see that in that equilibrium we have $f'(k) = r$, where r is the interest rate, or equivalently, the rate of return on capital.

This means that the condition for dynamic efficiency, which holds in the NGM, can be interpreted as the $r > g$ condition made famous by Piketty (2014) in his influential *Capital in the 21st Century*. The condition $r > g$ is what Piketty calls the "Fundamental Force for Divergence": an interest rate that exceeds the growth rate of the economy. In short, he argues that, if $r > g$ holds, then there will be a tendency for inequality to explode as the returns to capital accumulate faster than overall income grows. In Piketty's words:

> 'This fundamental inequality (...) will play a crucial role in this book. In a sense, it sums up the overall logic of my conclusions. When the rate of return on capital significantly exceeds the growth rate of the economy (...), then it logically follows that inherited wealth grows faster than output and income.' (pp. 25–26)

Does that mean that, were we to explicitly consider inequality in a context akin to the NGM we would predict it to explode along the BGP? Not so fast. First of all, when taking the model to the data, we could ask what k is. In particular, k can have a lot of human capital i.e. be the return to labour mostly, and this may help undo the result. In fact, it could even turn it upside down if human capital is most of the capital and is evenly distributed in the population. You may also want to see Acemoglu and Robinson (2015), who have a thorough discussion of this prediction. In particular, they argue that, in a model with workers and capitalists, modest amounts of social mobility – understood as a probability that some capitalists may become workers, and vice-versa – will counteract that force for divergence.

Yet the issue has been such a hot topic in the policy debate that two more comments on this issue are due.

First, let's understand better the determinants of labour and income shares. Consider a typical Cobb-Douglas production function:

$$Y = AL^{\alpha}K^{1-\alpha}. \tag{3.19}$$

With competitive factor markets, the FOC for profit maximisation would give:

$$w = \alpha AL^{\alpha-1}K^{1-\alpha}. \tag{3.20}$$

Computing the labour share using the equilibrium wage gives:

$$\frac{wL}{Y} = \frac{\alpha AL^{\alpha-1}K^{1-\alpha}L}{AL^{\alpha}K^{1-\alpha}} = \alpha, \tag{3.21}$$

which implies that for a Cobb-Douglas specification, labour and capital shares are constant. More generally, if the production function is

$$Y = \left(\beta K^{\frac{\varepsilon-1}{\varepsilon}} + \alpha\left(AL\right)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}} \text{ with } \varepsilon \in [0, \infty), \tag{3.22}$$

then ε is the (constant) elasticity of substitution between physical capital and labour. Note that when $\varepsilon \to \infty$, the production function is linear (K and L are perfect substitutes), and one can show that when $\varepsilon \to 0$ the production function approaches the Leontief technology of fixed proportions, in which one factor cannot be substituted by the other at all.

From the FOC of profit maximisation we obtain:

$$w = \left(\beta K^{\frac{\varepsilon-1}{\varepsilon}} + \alpha\left(AL\right)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{1}{\varepsilon-1}} \alpha A\left(AL\right)^{-\frac{1}{\varepsilon}}, \tag{3.23}$$

the labour share is now:

$$\frac{wL}{Y} = \frac{\alpha\left(\beta K^{\frac{\varepsilon-1}{\varepsilon}} + \alpha\left(AL\right)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{1}{\varepsilon-1}} A^{\frac{\varepsilon-1}{\varepsilon}} L^{-\frac{1}{\varepsilon}} L}{\left(\beta K^{\frac{\varepsilon-1}{\varepsilon}} + \alpha\left(AL\right)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}} = \alpha\left(\frac{AL}{Y}\right)^{\frac{\varepsilon-1}{\varepsilon}}. \tag{3.24}$$

Notice that as $\frac{L}{Y} \longrightarrow 0$, several things can happen to the labour share, and what happens depends on A and ε :

$$\text{If } \varepsilon > 1 \Longrightarrow \alpha\left(\frac{AL}{Y}\right)^{\frac{\varepsilon-1}{\varepsilon}} \longrightarrow 0 \tag{3.25}$$

$$\text{If } \varepsilon < 1 \Longrightarrow \alpha\left(\frac{AL}{Y}\right)^{\frac{\varepsilon-1}{\varepsilon}} \text{ increases.} \tag{3.26}$$