Lecture notes for Macro 475, Fall 2001

Per Krusell

Please do NOT distribute without permission

Comments and suggestions are welcome!

1 Introduction

These lecture notes cover a one-semester course. The overriding goal of the course is to begin provide methodological tools for advanced research in macroeconomics. The emphasis is on theory, although data guides the theoretical explorations. We build entirely on models with microfoundations, i.e., models where behavior is *derived* from basic assumptions on consumers' preferences, production technologies, information, and so on. Behavior is always assumed to be rational: given the restrictions imposed by the primitives, all actors in the economic models are assumed to maximize their objectives.

Macroeconomic studies emphasize decisions with a time dimension, such as various forms of investments. Moreover, it is often useful to assume that the time horizon is infinite. This makes dynamic optimization a necessary part of the tools we need to cover, and the first significant fraction of the course goes through, in turn, sequential maximization and dynamic programming. We assume throughout that time is discrete, since it leads to simpler and more intuitive mathematics.

The baseline macroeconomic model we use is based on the assumption of perfect competition. Current research often departs from this assumption in various ways, but it is important to understand the baseline in order to fully understand the extensions. Therefore, we will also spend significant time on the concepts of dynamic competitive equilibrium, both expressed in sequence form and recursively (using dynamic programming). In this context, it is important to study the welfare properties of our dynamic equilibria.

Infinite-horizon models can employ different assumptions about the time horizon of each economic actor. We study two extreme cases: (i) all consumers (really, dynasties) live forever—the infinitely-lived agent model—and (ii) consumers have finite and deterministic lifetimes but there are consumers of different generations alive at any point in time—the overlapping-generations model. These two cases share many features but also have important differences. Most of the course material will build on infinitely-lived agents, but we also study the overlapping-generations model in some depth.

Finally, many macroeconomic issues involve uncertainty. Therefore, we spend some time on how to introduce it into our models, both mathematically and in terms of economic concepts. Dynamic optimization, dynamic competitive analysis, dynamic welfare analysis, and uncertainty constitute the main topics for the first part of the course. A midterm exam should take place after these topics are covered.

We spend the second part of the course going over some important macroeconomic topics. These involve growth and business cycle analysis, asset pricing, fiscal policy, monetary economics, unemployment, and inequality. Here, few new tools are introduced; we will instead simply apply the tools from the first part of the course.

2 Motivation: Solow's growth model

Most modern dynamic models of macroeconomics build on the framework described in Solow's (1956) paper.[1] To motivate what is to follow, we start with a brief description of the Solow model. This model was set up to study a closed economy, and we will assume that there is a constant population.

2.1 The model

The model consists of some simple equations:

$$C_t + I_t = Y_t = F\left(K_t, L\right) \tag{1}$$

$$I_t = K_{t+1} - (1-\delta)\, K_t \tag{2}$$

$$I_t = sF\left(K_t, L\right) \tag{3}$$

and

$$\delta \in [0,1]\,,\, s \in [0,1]\,.$$

[1] No attempt is made here to properly assign credit to the inventors of each model. For example, the Solow model could also be called the Swan model, although usually it is not.

The equations in (1) are an accounting equation, saying that total resources are either consumed or invested, and an equation saying that total resources are given by the output of a production function with capital and labor as inputs. We take labor input to be constant at this point, whereas the other variables are allowed to vary over time. The accounting equation can also be interpreted in terms of technology: this is a one-good, or one-sector, economy, where the one good can be used both for consumption and as capital (investment). Equation (2) describes capital accumulation: the output good, in the form of investment, is used to accumulate the capital input, and capital depreciates geometrically: a constant fraction δ disintegrates every period.

Equation (3) is a behavioral equation. Unlike in the rest of the course, behavior is here assumed directly: a constant fraction of output is saved, independently of what the level of output is.

These equations together form a complete dynamic system—an equation system defining how its variables evolve over time—for some given F. That is, we know, in principle, what $\{K_t\}_{t=0}^{\infty}$ and $\{Y_t, C_t, I_t\}_{t=0}^{\infty}$ will be given any initial capital value K_0.

In order to analyze the dynamics, we now make some assumptions.

- $F(0, L) = 0$.
- $F_K(0, L) > \dfrac{\delta}{s}$.
- $\lim_{k \to \infty} sF_K(K, L) + (1 - \delta) < 1$.
- F is strictly concave in K and strictly increasing in K.

An example of a function satisfying these assumptions, and that will be used repeatedly in the course, is $F(K, L) = AK^{\alpha}L^{1-\alpha}$ with $0 < \alpha < 1$. This production function is called Cobb-Douglas; A is a productivity parameter, and α and $1 - \alpha$ denote the capital and labor share, respectively. Why they are called shares will be the subject of discussion later on.

The law of motion equation for capital may be rewritten:

$$K_{t+1} = (1 - \delta) K_t + sF(K_t, L).$$

Mapping K_t into K_{t+1} graphically, this can be pictured as in Figure **??**.

The intersection of the 45^oline with the savings function determines the stationary point. It can be verified that the system exhibits "global convergence" to the unique strictly positive steady state, K^*, that satisfies:

$$\begin{aligned} K^* &= (1 - \delta) K^* + sF(K^*, L) \\ \delta K^* &= sF(K^*, L) \text{ (there is a unique positive solution).} \end{aligned}$$

Given this information, we thus have

Theorem 1 $\exists K^* > 0 : \forall K_0 > 0,\ K_t \to K^*$.

Proof outline.

(1) Find a K^* candidate; show it is unique.

(2) If $K_0 > K^*$, show that $K^* < K_{t+1} < K_t \ \ \forall t \geq 0$ (using $K_{t+1} - K_t = sF(K_t, L) - \delta K_t$). If $K_0 < K^*$, show that $K^* > K_{t+1} > K_t \ \forall t > 0$.

(3) We have concluded that K_t is a monotone sequence, and that it is also bounded. Now use a math theorem: a monotone bounded sequence has a limit.

■

The proof of this theorem establishes not only global convergence but also that convergence is monotonic. The result is rather special in that it holds only under quite restrictive circumstances (for example, a one-sector model is a key part of the restriction).

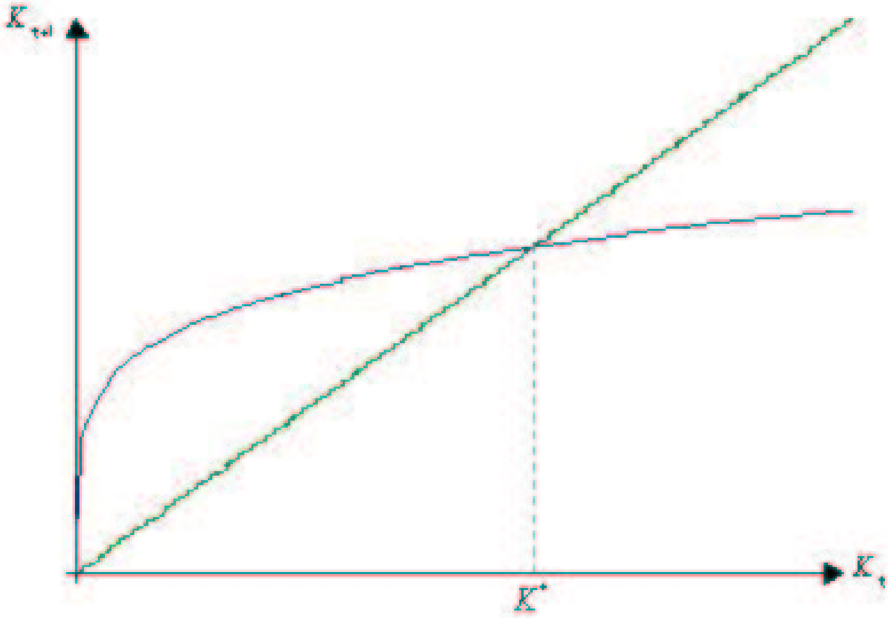

Figure 1: Convergence in the Solow model

2.2 Its uses

Growth. The Solow growth model is a corner block of many of the main setups in modern macroeconomic analysis. Its first and main use is that of understanding why output grows in the long run and what forms that growth takes. We will spend considerable time with that topic later. This will involve discussing what features of the production technology are important for long-run growth and it will involve analyzing the endogenous determination of productivity in a technological sense.

Consider, for example, the simple Cobb-Douglas case. In that case, α—the capital share—determines the shape of the law of motion function for capital accumulation. If α is close to one, the law of motion is close to linear in capital; if it is close to zero (but not exactly zero, the law of motion is quite nonlinear in capital. In terms of Figure 1, an α close to zero will make the steady state lower, and the convergence to the steady state will be quite rapid: from a given initial capital stock, few periods are necessary to get close to the steady state. If, on the other hand, α is close to one, the steady state is far to the right in the figure, and convergence will be slow.

When the production function is linear in capital—when α equals one—we have no positive steady state.[2] Suppose that $sA + 1 - \delta$ exceeds one. Then over time output would keep growing, and it would grow at precisely rate $sA + 1 - \delta$. Output and consumption would grow at that rate too. The "Ak" production technology is the simplest technology allowing "endogenous growth". That is the growth rate in the model is nontrivially determined, at least in the sense that it matters what behavior is for what the growth rate is. In the example, in particular, the savings rate is an important determinant of the growth rate. Savings rates that are very low will even make the economy shrink—if $sA + 1 - \delta$ goes below one. Keeping in mind that savings rates are probably influenced by government policy, such as taxation, this means that there would be a *choice*, both by individuals and government, of whether or not to grow.

The Ak model of growth emphasizes physical capital accumulation as the motor of prosperity. It is not the only way to think about growth, however. For example, one could model A more carefully and be specific about how productivity is enhanced over time via explicit decisions to accumulate R&D capital or human

[2]This statement is true unless $sA + 1 - \delta$ happened to equal 1.

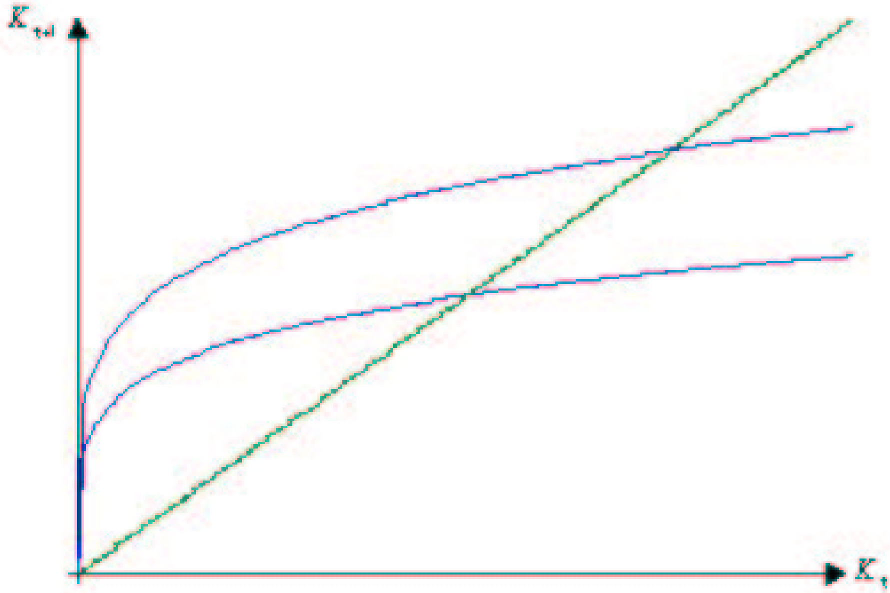

Figure 2: Random productivity in the Solow model

capital—learning. We will return to these different alternatives below.

In the context of understanding the growth of output, Solow also developed the methodology of "growth accounting", which is a way of breaking down the total growth of an economy into components: input growth and technology growth. We will discuss this later too; growth accounting remains a central tool for analyzing output and productivity growth over time and also for understanding differences between different economies in the cross-section.

Business Cycles. Many modern studies of business cycles also rely fundamentally of the Solow model. This includes real as well as monetary models. How can Solow's framework turn into a business cycle setup? Assume that the production technology will exhibit a stochastic component affecting the productivity of factors. For example, assume it is of the form

$$F = A_t \cdot \hat{F}(K_t, L)$$

where A_t is stochastic, for instance taking on two values: A_H, A_L. Retaining the assumption that savings rates are constant, we have what is depicted in Figure 2.

It is clear from studying this graph that as productivity realizations are high or low, output and total savings fluctuate. Will there be convergence to a steady state? In the sense of constancy of capital and other variables, steady states will clearly not be feasible here. However, another aspect of the convergence in deterministic model is inherited here: over time, initial conditions (the initial capital stock) lose influence and eventually—"after an infinite number of time periods"—the stochastic process for the endogenous variables will settle down and become stationary. Stationarity here is a statistical term, one that we will not develop in great detail in this course, although we will define it and use it for much simpler stochastic processes in the context of asset pricing. One element of stationarity in this case is that there will be a smallest compact set of capital stocks such that, once the capital stock is in this set, it never leaves the set: the "ergodic set". In the figure, this set is determined by the two intersections with the 45^oline.

Other topics. In other macroeconomic topics, such as monetary economics, labor, fiscal policy, and asset pricing, the Solow model is also commonly used. Then, other aspects need to be added to the framework, but Solow's one-sector approach is still very useful for talking about the macroeconomic aggregates.

2.3 Where next?

The model presented has the problem of relying on an exogenously determined savings rate. We saw that the savings rate, in particular, did not depend on the level of capital or output, nor on the productivity level. As stated in the introduction, this course aims to develop microfoundations. We would therefore like the savings behavior to be an outcome rather than an input into the model. To this end, the following chapters will introduce decision-making consumers into our economy. We will first cover decision making with a finite time horizon and then decision making when the time horizon is infinite. The decision problems will be phrased generally as well as applied to the Solow growth environment and other environments that will be of interest later.

3 Dynamic optimization

There are two common approaches to modeling real-life individuals: (i) they live a finite number of periods and (ii) they live forever. The latter is the most common apprach, but the former requires less mathematical sophistication in the decision problem. We will start with finite-life models and then consider infinite horizons.

We will also study two alternative ways of solving dynamic optimization problems: using sequential methods and using recursive methods. Sequential methods involve maximizing over sequences. Recursive methods—also labeled dynamic programming methods—involve functional equations. We begin with sequential methods and then move to recursive methods.

3.1 Sequential methods

3.1.1 A finite horizon

Consider a consumer having to decide a consumption stream for T periods. The preference ordering of the consumer of the consumption streams can be represented with the utility function

$$U\left(c_1, c_2, ..., c_T\right) \tag{4}$$

A standard assumption is that this function exhibits "additive separability", with stationary discounting weights.

$$U\left(c_1, c_2, ..., c_T\right) = \sum_{t=1}^{T} \beta^t u\left(c_t\right)$$

Notice that the per period (or instantaneous) utility index $u\left(\cdot\right)$ does not depend on time. Nevertheless, if instead we had $u_t\left(\cdot\right)$ the utility function $U\left(c_1, c_2, ..., c_T\right)$ would still be additively separable.

The powers of β are the discounting weights. They are called stationary because the ratio between the weights in any two different dates $t = i$ and $t = j > i$ only depends on the number of periods elapsed between i and j, and not on the actual value of i or j themselves.

The standard assumption is $0 < \beta < 1$, which corresponds to the observations that human beings seem to deem consumption at an early time more valuable than consumption further off in the future.

We now state the dynamic optimization problem associated with the neoclassical growth model in finite time.

$$\begin{array}{lll} \max\limits_{\{c_t,k_{t+1}\}_{t=0}^{T}} & \sum\limits_{t=0}^{T} \beta^t u(c_t) & \\ s.t. & c_t + k_{t+1} \leq f(k_t) = F(k_t, N) + (1-\delta) k_t & \\ & c_t \geq 0 \; t = 1, ..., T & \\ & k_{t+1} \geq 0 \; t = 1, ..., T & \\ & k_0 > 0 \text{ given} & \end{array} \tag{5}$$

This is a consumption-savings decision problem. It is, in this case, a "planning problem": there is no market where the individual might obtain an interest income from his savings, but rather savings yield production following the transformation rule $f(k_t)$.

The assumptions we will make on the production technology are the same as before. With respect to u, we will assume that it is strictly increasing. What's the implication of this? Notice that our resource constraint $c_t + k_{t+1} \leq f(k_t)$ allows for throwing goods away, since strict inequality is allowed. But the assumption that u is strictly increasing will imply that goods will not actually be thrown away, because they are valuable. We know in advance that the resource constraint will need to bind in our solution to this problem.

The solution method we will employ is straight out of standard optimization theory for finite-dimensional problems. In particular, we will make ample use of Kuhn-Tucker's theorem. The Kuhn-Tucker conditions:

(i) are necessary at an optimum (therefore, if we find a unique solution to the Kuhn-Tucker conditions, we are done!), provided a constraint qualification is met (we do not worry about it here);

(ii) are sufficient **if** the objective function is concave in the choice vector and the constraint set is convex.

We now characterize the solution further. It is useful to assume the following: $\lim\limits_{c \to 0} u'(c) = \infty$. This implies that $c_t = 0$ at any t cannot be optimal, so we can ignore the non-negativity constraint on consumption: we know in advance that it will not bind in our solution to this problem.

We write down the Lagrangian function:

$$L = \sum_{t=0}^{T} \beta^t \left[u(c_t) - \lambda_t \left[c_t + k_{t+1} - f(k_t) \right] + \mu_t k_{t+1} \right], \tag{[A]}$$

where we introduced the Lagrange/Kuhn-Tucker multipliers $\beta^t \lambda_t$ and $\beta^t \mu_t$ for our constraints.

The next step involves taking derivatives with respect to the decision variables c_t and k_{t+1} and stating the complete Kuhn-Tucker conditions. Before proceeding, however, let us take a look at an alternative formulation for this problem:

$$L = \sum_{t=0}^{T} \beta^t \left[u\left[f(k_t) - k_{t+1} \right] + \mu_t k_{t+1} \right] \tag{[B]}$$

Notice that we have made use of our knowledge of the fact that the resource constraint will be binding in our solution to get rid of the multiplier $\beta^t \lambda_t$. The two formulations are equivalent under the stated assumption on u. However, eliminating the multiplier $\beta^t \lambda_t$ might simplify the algebra. The multiplier may sometimes prove an efficient way of condensing information at the time of actually working out the solution.

We now solve the problem using [A]. The first-order conditions are:

$$\frac{\partial L}{\partial c_t} : \beta^t \left[u'(c_t) - \lambda_t \right] = 0 \qquad t = 0, ..., T$$

$$\frac{\partial L}{\partial k_{t+1}} : -\beta^t \lambda_t + \beta^t \mu_t + \beta^{t+1} \lambda_{t+1} f'(k_{t+1}) = 0 \qquad t = 0, ..., T-1$$

For period T,

$$\frac{\partial L}{\partial k_{T+1}} : -\beta^T \lambda_T + \beta^T \mu_T = 0.$$

The first-order condition under formulation [B] are:

$$\frac{\partial L}{\partial k_{t+1}} \quad : \quad -\beta^t u'(c_t) + \beta^t \mu_t + \beta^{t+1} u'(c_{t+1}) f'(k_{t+1}) = 0$$
$$\frac{\partial L}{\partial k_{T+1}} \quad : \quad -\beta^T u'(c_T) + \beta^T \mu_T = 0.$$

Finally, the Kuhn-Tucker conditions also include

$$\mu_t k_{t+1} = 0 t = 1, ..., T$$

$$\lambda_t \geq 0 t = 1, ..., T$$

$$k_{t+1} \geq 0 t = 1, ..., T$$

$$\mu_t \geq 0 t = 1, ..., T.$$

These conditions (the first of which is usually referred to as the complementary slackness condition) are the same for [A] and for [B].

The solution to both formulations become identical if in [A] we use $u'(c_t)$ to replace λ_t in the derivative $\dfrac{\partial L}{\partial k_{t+1}}$.

Now noting that $u'(c) > 0 \forall c$, we conclude that $\mu_T > 0$ in particular. This comes from the derivative of the Lagrangian with respect to k_{T+1}:

$$-\beta^T u'(c_T) + \beta^T \mu_T = 0.$$

But then this implies that $k_{T+1} = 0$: the consumer leaves no capital for after the last period, since he receives no utility from that capital and would rather use it for consumption during his lifetime. Of course, this is a trivial result, but its derivation is useful and will have an infinite-horizon counterpart that is less trivial.

The summary statement of the first-order conditions is then the "Euler equation":

$$\begin{aligned}
&u'[f(k_t) - k_{t+1}] = \beta u'[f(k_{t+1}) - k_{t+2}] f'(k_{t+1}) \quad t = 0, ..., T-1 \\
&k_0 \text{ given} \\
&k_{T+1} = 0
\end{aligned}$$

where the capital sequence is what we need to solve for. The Euler equation is sometimes referred to as a "variational" condition (as part of "calculus of variation"): given to boundary conditions k_t and k_{t+2}, it represents the idea of varying the intermediate value k_{t+1} so as to achieve the best outcome. Piecing together these variational conditions, therefore, we notice that there are a total of $T+2$ equations and $T+2$ unknowns—the unknowns are a sequence of capital stocks with an initial and a terminal condition. This is called a *difference equation* in the capital sequence. It is a *second-order* difference equation because there are two lags of capital in the equation. Since the number of unknowns is equal to the number of equations, the difference equation system will typically have a solution, and under appropriate assumptions on primitives, there will be only one such solution. We will now briefly look at conditions under which there is only one solution to the first-order conditions or, alternatively, under which the first-order conditions are sufficient.

What we need to assume is that u is concave. Using formulation [A], then we know that $U = \sum_{t=0}^{T} u(c_t)$ is concave in the vector $\{c_t\}$, since the sum of concave functions is concave. Moreover, the constraint set is convex in $\{c_t, k_{t+1}\}$, given that we assume concavity of f (this is straightforward to check using the definition of a convex set and of a concave function). So concavity of the functions u and f make the overall objective concave and the choice set convex, and thus the first-order conditions are sufficient. Alternatively, using formulation [B], since $u(f(k_t) - k_{t+1})$ is concave in (k_t, k_{t+1}), which follows from the fact that u is concave and increasing and that f is concave, the objective is concave in $\{k_{t+1}\}$. The constraint set in formulation [B] is clearly convex, since all it requires is $k_{t+1} \geq 0$ for all t.

Finally, a unique solution (to the problem as such as well as to the first-order conditions) is obtained if the objective is strictly concave, which we have if u is strictly concave.

To interpret the key equation for optimization, the Euler equation, it is useful to break it down in three components:

$$\underbrace{u'(c_t)}_{\substack{\text{Utility lost if you} \\ \text{invest "one" more} \\ \text{unit } \to \text{ marginal} \\ \text{cost of saving}}} = \underbrace{\beta u'(c_{t+1})}_{\substack{\text{Utility increase} \\ \text{next period per} \\ \text{unit of increased } c_{t+1}}} \cdot \underbrace{f'(k_{t+1})}_{\substack{\text{Return on the} \\ \text{invested unit: by how} \\ \text{many units next period's} \\ c \text{ can increase}}}$$

Thus, because of the concavity of u, equalizing the marginal cost of saving to the marginal benefit of saving is a condition for an optimum.

How do primitives affect savings behavior? We can identify three component determinants of saving: the concavity of utility, the discounting, and the return to saving. Their effects are described in turn.

(i) Consumption "smoothing": if the utility function is *strictly* concave, the individual prefers a smooth consumption stream.

$$\textit{Example:} \beta f'(k_{t+1}) = 1 \;\Rightarrow u'(c_t) = u'(c_{t+1}) \underbrace{\qquad \Rightarrow \qquad}_{\text{if } u \text{ is strictly concave}} c_t = c_{t+1}$$

(ii) Impatience: via β, we see that a low β (a low discount factor, or a high discount rate $\frac{1}{\beta} - 1$) $\Rightarrow$ will tend to be associated with low c_{t+1}s and high c_ts.

(iii) The return on savings: $f'(k_{t+1})$ clearly also affects behavior, but its effect on consumption cannot be signed unless we make more specific assumptions. Moreover, k_{t+1} is endogenous, so when f' nontrivially depends on it, we cannot vary the return independently. The case when f' is a constant, such as in the Ak growth model, is more convenient. We will return to it below.

To gain some more detailed understanding of the determinants of savings, let us study some examples.

Logarithmic utility. Let the felicity index be

$$u(c) = \log c$$

and the production technology be represented by the function

$$f(k) = R \cdot k.$$

Notice that this amounts to a linear function with exogenous marginal return R on investment.

The Euler equation becomes:

$$\begin{aligned} u'(c_t) &= \beta \cdot u'(c_{t+1}) \cdot \frac{f'(k_{t+1})}{R} \\ \frac{1}{c_t} &= \frac{\beta \cdot R}{c_{t+1}} \\ &\Rightarrow c_{t+1} = \beta \cdot R \cdot c_t. \end{aligned} \tag{EE}$$

The optimal path has consumption growing at the rate βR, and it is constant between any two periods. From the resource constraint (recall that it binds):

$$\begin{cases} c_0 + k_1 = Rk_0 \\ c_1 + k_2 = Rk_1 \\ \vdots \\ c_T + k_{T+1} = Rk_T \\ k_{T+1} = 0 \end{cases}$$

With repeated substitutions, we obtain the "consolidated" or "intertemporal" budget constraint:

$$c_0 + \frac{1}{R}c_1 + \frac{1}{R^2}c_2 + ... + \frac{1}{R^T}c_T = Rk_0$$

The left-hand side is the present value of the consumption stream, and the right hand side is the present value of income. Using the optimal consumption growth rule $c_{t+1} = \beta R \cdot c_t$,

$$\begin{aligned} c_0 + \frac{1}{R}\beta R \cdot c_0 + \frac{1}{R^2}\beta^2 R^2 \cdot c_0 + ... + \frac{1}{R^T}\beta^T R^T \cdot c_0 &= Rk_0 \\ c_0 \cdot \left[1 + \beta + \beta^2 + ... + \beta^T\right] &= Rk_0. \end{aligned}$$

This implies

$$\boxed{c_0 = \frac{Rk_0}{1 + \beta + \beta^2 + ... + \beta^T}}$$

We are now able to study the effects on the consumer's behavior of changes in the marginal return on savings, R. An increase in R will cause a rise in consumption *in all periods*. Crucial to this result is the chosen form for the utility function. Logarithmic utility has the property that income and substitution effects, when they go in opposite directions, offset each other. Changes in R have two components: a change in relative prices (of consumption in different periods) and a change in present-value income: $R \cdot k_0$. With logarithmic utility, a relative price change between two goods will make the consumption of the favored good go up whereas the consumption of other good will remain at the same level. The unfavored good will not be consumed in a lower amount since there is a positive income effect of the other good being cheaper, and that effect will be spread over both goods. Thus, the period 0 good will be unfavored in our example (since all other goods have lower price relative to good 0 if R goes up), and its consumption level will not decrease. The consumption of good 0 will in fact increase because total present-value income is multiplicative in R.

Next assume that the sequence of interest rates is not constant, but that instead $\{R_t\}_{t=0}^{T}$ with R_t different at each t. The consolidated budget constraint now reads:

$$c_0 + \frac{1}{R_1}c_1 + \frac{1}{R_1 R_2}c_2 + \frac{1}{R_1 R_2 R_3}c_3 + ... + \frac{1}{R_1 ... R_T}c_T = a_0 R_0.$$

Plugging in the optimal path rule from (EE), one obtains

$$c_0 \cdot \left[1 + \beta + \beta^2 + ... + \beta^T\right] = a_0 R_0$$

from where

$$\begin{aligned} c_0 &= \frac{a_0 R_0}{1 + \beta + \beta^2 + ... + \beta^T} \\ c_1 &= \frac{a_0 R_0 R_1 \beta}{1 + \beta + \beta^2 + ... + \beta^T} \\ &\vdots \\ c_t &= \frac{a_0 R_0 ... R_t \beta^t}{1 + \beta + \beta^2 + ... + \beta^T}. \end{aligned}$$

Now note the following comparative statics:

$$\begin{aligned} R_t \uparrow \quad &\Rightarrow c_0, c_1, ..., c_{t-1} \text{ are unaffected} \\ &\Rightarrow \text{ savings at } 0,\ ...,\ t-1 \text{ are unaffected} \end{aligned}$$

In the logarithmic utility case, if the return between t and $t+1$ changes, consumption and savings remain unaltered until $t-1$!!

A slightly more general utility function. Let us introduce the most common additively separable utility function in macroeconomics: the CES (constant elasticity of substitution) function:

$$u(c) = \frac{c^{1-\sigma} - 1}{1-\sigma}$$

This function has as special cases:

$\sigma = 0$	linear utility
$\sigma > 0$	strictly concave utility
$\sigma = 1$	logarithmic utility
$\sigma = \infty$	not possible, but this is usually referred to as Leontief utility function

Let us define the intertemporal elasticity of substitution (IES):

$$IES \equiv \frac{\frac{d\left(\frac{c_{t+k}}{c_t}\right)}{\frac{c_{t+k}}{c_t}}}{\frac{dR_{t,t+k}}{R_{t,t+k}}}$$

We will show that all the special cases of the CES function have constant intertemporal elasticity of substitution equal to $\frac{1}{\sigma}$. We begin with the Euler equation:

$$u'(c_t) = \beta u'(c_{t+1}) R_{t+1}.$$

Replacing repeatedly, we have

$$\begin{aligned}
u'(c_t) &= \beta^k u'(c_{t+k}) \underbrace{R_{t+1}R_{t+2}...R_{t+k}}_{\equiv R_{t,t+k}} \\
u'(c) &= c^{-\sigma} \Rightarrow c_t^{-\sigma} = \beta^k c_{t+k}^{-\sigma} R_{t,t+k} \\
\frac{c_{t+k}}{c_t} &= \left(\beta^k\right)^{\frac{1}{\sigma}} (R_{t,t+k})^{\frac{1}{\sigma}}.
\end{aligned}$$

This means that our elasticity measure becomes

$$\frac{\frac{d\left(\frac{c_{t+k}}{c_t}\right)}{\frac{c_{t+k}}{c_t}}}{\frac{dR_{t,t+k}}{R_{t,t+k}}} = \frac{d\log\frac{c_{t+k}}{c_t}}{d\log R_{t,t+k}} = \frac{1}{\sigma}.$$

When $\sigma = 1$, expenditure shares do not change: this is the logarithmic case. When $\sigma > 1$, an increase in $R_{t,t+k}$ would lead c_t to go up and savings to go down: the income effect, leading to smoothing across all goods, is larger than substitution effect. Finally, when $\sigma < 1$, the substitution effect is stronger: savings go up whenever $R_{t,t+k}$ goes up.

When $\sigma = 0$, the elasticity is infinite and savings respond discontinuously to $R_{t,t+k}$.

3.1.2 Infinite horizon

Why should macroeconomists study the case of an infinite time horizon? There are at least two reasons:

1. *Altruism*: People do not live forever, but they may care about their offspring. Let $u(c_t)$ denote the utility flow to generation t. We can then interpret β^t as the weight an individual attaches to the utility enjoyed by his descendants t generations down the family tree. His total joy is given by: $\sum_{t=0}^{\infty} \beta^t u(c_t)$. A $\beta < 1$ thus implies that the individual cares more about himself than about his descendants.

   With generational overlap, the utility function would look similar:

   $$\sum_{t=0}^{\infty} \beta^t \underbrace{\left[u(c_{yt}) + \delta u(c_{ot})\right]}_{\text{utility flow to generation } t}.$$

   The existence of bequests indicates that there is altruism. However, bequests can also be of an entirely selfish, precautionary nature: when the life-time is unknown, as it is in practice, bequests would then be accidental and simply reflect the remaining buffer the individual kept for the possible remainder of his life. An argument for why bequests may not be entirely accidental is that annuity markets are not used very much. Annuity markets allow you to effectively insure against living "too long", and would thus make bequests disappear: all your wealth would be put into annuities and disappear upon death.

   It is important to point out that the time horizon for an individual only becomes truly infinite if the altruism takes the form of caring about the utility of the descendants. If, instead, utility is derived from the act of giving itself, without reference to how the gift influences others' welfare, the individual's problem again becomes finite. Thus, if I live for one period and care about how much I give, my utility function might be $u(c) + v(b)$, where v measures how much I enjoy giving bequests, b. Although b subsequently shows up in another agent's budget and influences his choices and welfare, those effects are irrelevant for the decision of the present agent, and we have a simple static framework. This model is usually referred to as the "warm glow" model (the giver feels a warm glow from giving).

   For a variation, think of an individual (or a dynasty) that, if still alive, each period dies with probability π. Its expected lifetime utility from a consumption stream $\{c_t\}_{t=0}^{\infty}$ is then given by

   $$\sum_{t=0}^{\infty} \beta^t \pi^t u(c_t).$$

   This framework—the "perpetual-youth" model, or, perhaps better, the "sudden-death" model—is sometimes used in applied contexts. Analytically, it looks like the infinite-life model, only with the difference that the discount factor is $\beta\pi$. These models are thus the same on the individual level. On the aggregate level, they are not, since the sudden-death model carries with it the assumption that a deceased dynasty is replaced with a new one: it is, formally speaking, an overlapping-generations model (see more on this below), and as such it is different in certain key respects.

   Finally, one can also study explicit games between players of different generations. We may assume that parents care about their children, that sons care about their parents as well, and that each of their activities in part is motivated by this altruism, leading to intervivos gifts as well as bequests. Since such models lead us into game theory rather quickly, and therefore typically to more complicated characterizations, we will assume that altruism is unidirectional.

2. *Simplicity*: Many macroeconomic models with a long time horizon tend to show very similar results to infinite-horizon models if the horizon is long enough. Infinite-horizon models are stationary in nature—the remaining time horizon does not change as we move forward in time—and their characterization can therefore often be obtained more easily than when the time horizon changes over time.

   The similarity in results between long- and infinite-horizon setups is is not present in all models in economics. For example, in dynamic game theory the Folk Theorem means that the extension from a long (but finite) to an infinite horizon introduces a qualitative change in the model results. The typical example of this "discontinuity at infinity" is the prisoner's dilemma repeated a finite number of times, leading to a unique, noncooperative outcome, versus the same game repeated an infinite number of times, where there is a large set of equilibria.

Models with an infinite time horizon demand more advanced mathematical tools. Consumers in our models are now choosing infinite sequences. These are no longer elements of Euclidean space $\Re^n$, which was used for our finite-horizon case. A basic question is when solutions to a given problem exist. Suppose we are seeking to maximize a function $U(x)$, $x \in S$. If $U(\cdot)$ is a continuous function, then we can invoke Weierstrass's theorem provided that the set S meets the appropriate conditions: S needs to be nonempty and compact. For $S \subset \Re^n$, compactness simply means closedness and boundedness. In the case of finite horizon, recall that x was a consumption vector of the form $(c_1, ..., c_T)$ from a subset S of $\Re^T$. In these cases, it was usually easy to check compactness. But now we have to deal with larger spaces; we are dealing with infinite-dimensional sequences $\{k_t\}_{t=0}^{\infty}$. Several issues arise. How do we define continuity in this setup? What is an open set? What does compactness mean? We will not answer these questions here, but we will bring up some specific examples of situations when maximization problems are ill-defined, that is, when they have no solution.

A. Examples where utility may be unbounded

Continuity of the objective requires boundedness. When will U be bounded? If two consumption streams yield "infinite" utility, it is not clear how to compare them. The device chosen to represent preference rankings over consumption streams is thus failing. But is it possible to get unbounded utility? How can we avoid this pitfall?

Utility may become unbounded for one of many reasons. Although these reasons interact, let us consider each one independently.

Preference requirements: Consider a plan specifying equal amounts of consumption goods for each period, throughout eternity:

$$\{c_t\}_{t=0}^{\infty} = \{\bar{c}\}_{t=0}^{\infty}$$

Then the value of this consumption stream according to the chosen time-separable utility function representation is computed by:

$$U = \sum_{t=0}^{\infty} \beta^t u(c_t) = \sum_{t=0}^{\infty} \beta^t u(\bar{c})$$

What is a necessary condition for U to take on a finite value in this case? The answer is $\beta < 1$: under this parameter specification, the series $\sum_{t=0}^{\infty} \beta^t u(\bar{c})$ is convergent, and has a finite limit. If $u(\cdot)$ has the IES parametric form, then the answer to the question of convergence will involve not only β, but also σ.

Alternatively, consider a constantly increasing consumption stream:

$$\{c_t\}_{t=0}^{\infty} = \left\{c_0 \cdot (1+\gamma)^t\right\}_{t=0}^{\infty}$$

Is $U = \sum_{t=0}^{\infty} \beta^t u(c_t) = \sum_{t=0}^{\infty} \beta^t u\left(c_0 \cdot (1+\gamma)^t\right)$ bounded? Notice that the argument in the instantaneous utility index $u(\cdot)$ is increasing without bound, while for $\beta < 1$ β^t is decreasing to 0. This seems to hint that the key to having a convergent series this time lies in the form of $u(\cdot)$ and in how it "processes" the increase in the value of its argument. In the case of IES utility representation, the the relationship between β, σ, and γ is thus the key to boundedness. In particular, boundedness requires $\beta\gamma^{1-\sigma} < 1$. If $\beta < 1$, the boundedness fails if $\sigma < 1$ and $\gamma > 1$: utility will be $+\infty$ in this case, so if it is feasible to choose a long-run consumption growth at this rate γ, utility would have no upper bound. Still with $\beta < 1$, the boundedness condition also fails if $\sigma > 1$ and $\gamma < 1$; in this case, utility will be $-\infty$. What this might mean is that if it is not feasible to make consumption grow faster than γ (so that it is actually shrinking), all consumption paths are equally horrible, yielding so low a utility outcome that the maximization problem is ill-defined. Finally, note that $\beta > 1$ is possible to allow, since the utility function would be bounded in this case, provided for example that consumption is growing ($\gamma > 1$) with a large enough σ. Then utility would be a bounded (negative) number. This case is curious, however, as it would mean that another feasible path, one where consumption is constant, would be ill-defined ($-\infty$), thus making any constant consumption paths uncomparable, even though one would think that a higher constant consumption is always better than a lower constant consumption.[3]

Two other issues are involved in the question of boundedness of utility. One is technological, and the other may be called institutional.

A technological consideration: Technological restrictions are obviously necessary in some cases, as illustrated indirectly above. Let the technological constraints facing the consumer be represented by the budget constraint:

$$\begin{aligned} c_t + k_{t+1} &= R \cdot k_t \\ k_t &\geq 0 \end{aligned}$$

This constraint needs to hold for all time periods t (this is just the "Ak" case already mentioned). This implies that consumption can grow by (at most) a rate of R. A given rate R may thus be so high that it leads to unbounded utility, as shown above.

The institutional framework: Some things simply cannot happen in an organized society. One of these is so dear to analysts modelling infinite-horizon economies that it has a name of its own. It expresses the fact that if an individual announces that he plans to borrow and never pay back, then he will not be able to find a lender. The requirement that "no Ponzi games are allowed" thus represents this institutional assumption, and it sometimes needs to be added formally to the budget constraints of a consumer.

First, let $c_t \leq \bar{c}\ \forall t \geq 1$; that is, let consumption be bounded for every t, but let the initial consumption amount, c_0, be arbitrary. Suppose we endow a consumer with a given initial amount of net assets, a_0. These represent (real) claims against other agents. The constraint set is assumed to be

$$\begin{aligned} c_t + a_{t+1} &= R \cdot a_t t = 0, 1, ..., ... \\ c_t &= \bar{c}\ t = 1, 2, ..., ... \end{aligned}$$

If $a_1 \equiv -c_0 + a_0 R < 0$, our agent is thus borrowing. As for subsequent periods,

$$\begin{aligned} a_2 &= a_1 R - \bar{c} = a_0 R^2 - c_0 R - \bar{c} \\ &\vdots \\ a_t &= R^t \cdot \left[a_0 - \frac{c_0}{R} \right] - \bar{c} \frac{R^{t-1} - 1}{R - 1}. \end{aligned}$$

Thus, for a given initial wealth, a_0, if the initial consumption level, c_0, or the long-run one, $\bar{c}$, are too high, the consumer's debt will be growing without bound at rate R, and it will never be repaid. We rule out this situation by imposing the no-Ponzi-game (nPg) condition, by explicitly adding the restriction that:

$$\lim_{t \to \infty} \frac{a_t}{R^t} \geq 0.$$

Intuitively, this means that in present-value terms, the agent is cannot engage in borrowing and lending so that his "terminal position" is negative, since this means that he would borrow and not pay back.

Can we use the nPg condition to simplify, or "consolidate", the sequence of budget constraints? By repeatedly replacing T times, we obtain

$$\sum_{t=0}^{T} c_t \cdot \frac{1}{R^t} + \frac{a_T}{R^{T+1}} \leq a_0 \cdot R.$$

[3]The curiosity, which seems to contradict that the preferences satisfy a basic monotonicity axiom, is partially resolved by allowing the two constant paths to be compared with a so-called "overtaking criterion": one path is preferred to another if there is a time horizon such that, for any finite horizon longer than it, it gives higher utility. This criterion would also give the same answer as the traditional calculation when the utility of the infinite stream is bounded. There is a potential remaining problem if two paths are not comparable using this criterion (for example, with $\beta = 1$ and $\sigma = 0$, any two oscillating paths with the same mean are not comparable).

By the nPg condition, we then have

$$\lim_{T\to\infty}\left(\sum_{t=0}^{T} c_t \cdot \frac{1}{R^t} + \frac{a_T}{R^{T+1}}\right) = \lim_{T\to\infty}\sum_{t=0}^{T} c_t \cdot \frac{1}{R^t} \equiv \sum_{t=0}^{\infty} c_t \cdot \frac{1}{R^t},$$

and since the inequality is valid for every T,

$$\sum_{t=0}^{\infty} c_t \cdot \frac{1}{R^t} \leq a_0 \cdot R.$$

This is the consolidated budget constraint. In practice, we will often use a version of nPg with equality.

B. Sufficient conditions

Maximization of utility under an infinite horizon will in most parts involve the same mathematical techniques as in the finite-horizon case. In particular, we will make use of (Kuhn-Tucker) first-order conditions: barring corner constraints, we will choose a path such that the marginal effect of any choice variable on utility is zero. In particular, consider the sequences that the consumer chooses for his consumption and accumulation of capital. The first-order conditions will then lead to an Euler equation, which is defined for any path for capital, beginning with an initial value k_0. In the case of finite time horizon, it did not make sense for the agent to invest in the final period T, since no utility would be enjoyed from consuming goods at time $T+1$ when the economy is inactive. This final zero capital condition was key to determining the optimal path of capital: it provided us with a terminal condition for a difference equation system. In the case of infinite time horizon, there is no such final T: the economy will continue forever. Therefore, the difference equation that characterizes the first-order condition may have an infinite number of solutions. We will thus need some other way of pinning down the consumer's choice, and it turns out that the missing condition is analogous to the requirement that the capital stock be zero at $T+1$, for else the consumer could increase his utility. The missing condition, which we will now discuss in detail, is called the *transversality* condition. It is, typically, a necessary condtion for an optimum, and it expresses the following simple idea: it cannot be optimal for the consumer to choose a capital sequence such that, in present-value utility terms, the shadow value of k_t remains positive as t goes to infinity. This could not be optimal because it would represent saving too much: a reduction in saving would thus still be feasible and increase utility.

We will not prove the necessity of the transversality condition here. We will, however, provide a sufficiency condition. Suppose that we have a convex maximization problem (utility is concave and the constraint set convex) and a sequence $\{k_{t+1}\}_{t=1}^{\infty}$ satisfying the Kuhn-Tucker first-order conditions for a given k_0. Is $\{k_{t+1}\}_{t=1}^{\infty}$ a maximum? We did not formally prove a similar proposition in the finite-horizon case (we merely referred to math texts), but we will here, and the proof can also be used for finite-horizon setups.

Sequences satisfying the Euler equations but that do not maximize the programming problem come up quite often. We would like to have a systematic way of distinguishing between maxima and other critical points (in $\Re^{\infty}$) that are not the solution we are looking for. Fortunately, the transversality condition helps us here: if a sequence $\{k_{t+1}\}_{t=1}^{\infty}$ satisfies both the Euler equations and the transversality condition, then it maximizes the objective function. Formally, we have the following

Proposition 2 *Consider the programming problem*

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \cdot F(k_t,\ k_{t+1})$$
$$s.t. \quad k_{t+1} \geq 0 \forall t.$$

(An example is $F(x,y) = u[f(x) - y]$.)

If $\left\{k_{t+1}^*\right\}_{t=0}^{\infty}$, $\{\mu_t^*\}_{t=0}^{\infty}$ satisfy *(i)* $k^*_{t+1} \geq 0 \ \forall t$

(ii) Euler Equation: $F_2\left(k^*_t,\ k^*_{t+1}\right) + \beta F_1\left(k^*_{t+1},\ k^*_{t+2}\right) + \mu^*_t = 0 \ \forall t$

(iii) $\mu^*_t \geq 0,\ \mu^*_t \cdot k^*_{t+1} \geq 0 \ \forall t$

(iv) $\lim\limits_{T\to\infty} \beta^t F_1\left(k^*_t,\ k^*_{t+1}\right) k^*_t = 0$

and $F(x,y)$ *is concave in* (x,y) *and increasing in its first argument, then* $\left\{k^*_{t+1}\right\}^{\infty}_{t=0}$ *maximizes the objective.*

Proof. Consider any alternative feasible sequence $\mathbf{k} \equiv \{k_{t+1}\}^{\infty}_{t=0}$. Feasibility is tantamount to $k_{t+1} \geq 0$ $\forall t$. We want to show that for any such sequence,

$$\lim_{T\to\infty} \sum_{`t=0}^{T} \beta^t \left[F\left(k^*_t,\ k^*_{t+1}\right) - F\left(k_t,\ k_{t+1}\right)\right] \geq 0.$$

Define

$$A_T(\mathbf{k}) \equiv \sum_{t=0}^{T} \beta^t \left[F\left(k^*_t,\ k^*_{t+1}\right) - F\left(k_t,\ k_{t+1}\right)\right].$$

We will to show that, as T goes to infinity, $A_T(\mathbf{k})$ is bounded below by zero.

By concavity of F,

$$A_T \geq \sum_{t=0}^{T} \beta^t \left[F_1\left(k^*_t,\ k^*_{t+1}\right) \cdot \left(k^*_t - k_t\right) + F_2\left(k^*_t,\ k^*_{t+1}\right) \cdot \left(k^*_{t+1} - k_{t+1}\right)\right].$$

Now notice that for each t, k_{t+1} shows up twice in the summation. Hence we can rearrange the expression to read

$$\begin{aligned} A_T \quad \geq \quad & \sum_{t=0}^{T} \beta^t \left\{\left(k^*_{t+1} - k_{t+1}\right) \cdot \left[F_2\left(k^*_t,\ k^*_{t+1}\right) + \beta F_1\left(k^*_{t+1},\ k^*_{t+2}\right)\right]\right\} + \\ & + F_1\left(k^*_0,\ k^*_1\right) \cdot \left(k^*_0 - k_0\right) + \beta^T F_2\left(k^*_T,\ k^*_{T+1}\right) \cdot \left(k^*_{T+1} - k_{T+1}\right). \end{aligned}$$

Some information contained in the first-order conditions will now be useful:

$$F_2\left(k^*_t,\ k^*_{t+1}\right) + \beta F_1\left(k^*_{t+1},\ k^*_{t+2}\right) = -\mu^*_t,$$

together with $k^*_0 - k_0 = 0$ (k_0 can only take on one feasible value), allows us to derive

$$A_T \geq \sum_{t=0}^{T} \beta^t \mu^*_t \cdot \left(k_{t+1} - k^*_{t+1}\right) + \beta^T F_2\left(k^*_T,\ k^*_{T+1}\right) \cdot \left(k^*_{T+1} - k_{T+1}\right).$$

Next, we use the complementary slackness conditions and the implication of the Kuhn-Tucker conditions that

$$\mu^*_t k_{t+1} \quad \geq \quad 0$$

to conclude that $\mu^*_t \cdot \left(k_{t+1} - k^*_{t+1}\right) \geq 0$. In addition, $F_2\left(k^*_T,\ k^*_{T+1}\right) = -\beta F_1\left(k^*_{T+1},\ k^*_{T+2}\right) - \mu^*_{T+1}$, so we obtain

$$A_T \geq \sum_{t=0}^{T} \beta^t \mu^*_t \cdot \left(k_{t+1} - k^*_{t+1}\right) + \beta^T \left[\beta F_1\left(k^*_{T+1},\ k^*_{T+2}\right) + \mu^*_{T+1}\right] \cdot \left(k_{T+1} - k^*_{T+1}\right).$$

Since we know that $\mu^*_{t+1} \cdot \left(k_{t+1} - k^*_{t+1}\right) \geq 0$, the value of the summation will be reduced if we suppress nonnegative terms:

$$A_T \geq \beta^{T+1} F_1\left(k^*_{T+1},\ k^*_{T+2}\right) \cdot \left(k_{T+1} - k^*_{T+1}\right) \geq -\beta^{T+1} F_1\left(k^*_{T+1},\ k^*_{T+2}\right) \cdot k^*_{T+1}.$$

Finally, as T goes to infinity, the right-hand side of the last inequality goes to zero by the transversality. That is, we have shown that the utility implied by the candidate path must be higher than that implied by the alternative. ■

In the finite horizon case, k^*_{T+1} would have been the level of capital left out for the day after the (perfectly foreseen) end of the world; a requirement for an optimum in that case is clearly $k^*_{T+1} = 0$. In present-value utility terms, one might alternatively require $k^*_{T+1}\beta^T\lambda^*_T = 0$, where $\beta^t\lambda^*_t$ is the present-value utility evaluation of an additional unit of resources in period t. The transversality condition can be given this interpretation: F_{1t} is the marginal addition of utils in period t from increasing capital in that period, so the transversality condition simply says that the value (discounted into present-value utils) of each additional unit of capital at infinity times the actual amount of capital hs to be zero. If this expression were not met (we are now, incidentally, making a heuristic argument for necessity), it would pay off for the consumer to modify such a capital path and increase consumption for an overall increase in utility without violating feasibility.[4]

The no-Ponzi-game and the transversality conditions play very similar roles in dynamic optimization in a purely mechanical sense (at least if the nPg condition is interpreted with equality). In fact, they can typically be shown to be the same condition, if one also assumes that the first-order condition is satisfied. However, the two conditions are conceptually very different. The nPg condition is a restriction on the choices of the agent. In contrast, the transversality condition is a prescription for how to behave optimally *given* a choice set.

3.2 Dynamic Programming

The models we are concerned with consist of a more or less involved dynamic optimization problem and a resulting optimal consumption plan that solves it. Our approach up to now has been to look for a sequence of real numbers $\left\{k^*_{t+1}\right\}_{t=0}^{\infty}$ that generates an optimal consumption plan. In principle, this has involved searching for a solution to an infinite sequence of equations—a difference equation (the Euler equation). The search for a sequence is sometimes impractical, and not always intuitive. An alternative approach is often available, however, one which is useful conceptually as well as for computation (both analytical and, especially, numerical computation). It is called dynamic programming. We will now go over the basics of this approach. The focus will be on concepts, as opposed to on the mathematical aspects or on the formal proofs (these will be an important aspect of next semester's macro course).

Key to dynamic programming is to think of dynamic decisions as being made not once and for all but recursively: time period by time period. The savings between t and $t+1$ are thus decided on at t, and not at 0. A problem is *stationary* whenever the structure of the choice problem that a decision maker faces is identical at every point in time. As an illustration, in the examples that we have seen so far, we posited a consumer placed at the beginning of time choosing his infinite future consumption stream given an initial capital stock k_0. As a result, out came a sequence of real numbers $\left\{k^*_{t+1}\right\}_{t=0}^{\infty}$ indicating the level of capital that the agent will choose to hold in each period. But once he has chosen a capital path, suppose that we let the consumer abide to it for, say, T periods. At $t = T$ he will find then himself with the k^*_T decided on initially. If at that moment we told the consumer to forget about his initial plan and asked him to decide on his consumption stream again, from them onwards, using as new initial level of capital $k_0 = k^*_T$, what sequence of capital would he choose? If the problem is *stationary* then for any two periods $t \neq s$,

$$k_t = k_s \Rightarrow k_{t+j} = k_{s+j}$$

for all j larger than zero. That is, he would not change his mind if he could decide all over again.

This means that, if a problem is stationary, we can think of a function that, for every period t, assigns to each possible initial level of capital k_t an optimal level for next period's capital k_{t+1} (and therefore an optimal level of current period consumption): $k_{t+1} = g(k_t)$. Stationarity means that the function $g(\cdot)$ has no other argument than current capital. In particular, the function does not vary with time. We will refer to $g(\cdot)$ as a *decision rule.*

We have defined stationarity above in terms of decisions—in terms of properties of the solution to a dynamic problem. What types of dynamic problems are stationary? Intuitively, a dynamic problem is

[4]This necessity argument clearly requires that utility is strictly increasing in capital.

stationary if one can capture what is relevant information for the decision maker in a way that does not involve time. In our neoclassical growth framework, with a finite horizon, time is important, and the problem is not stationary: it matters how many periods are left—the decision problem changes character as time passes. With an infinite time horizon, however, the remaining horizon is the same at each point in time. The only changing feature of the consumer's problem in the infinite-horizon neooclassical growth economy is his initial capital stock; hence, his decisions will not depend on anything but this capital stock. Whatever the relevant information is for a consumer solving a dynamic problem, we will refer to it as his *state variable.* So the state variable for the planner in the one-sector neoclassical growth context is the current capital stock.

The heuristic information above can be expressed more formally as follows. The simple mathematical idea that $\max_{x,y} f(x,y) = \max_y\{\max_x f(x,y)\}$ (if each of the max operators are well-defined) allows us to maximize "in steps": first over x, given y, and then the remainder (where we can think of x as a function of y) over y. If we do this over time, the idea would be to maximize over $\{k_{s+1}\}_{s=t}\infty$ first by choice of $\{k_{s+1}\}_{s=t+1}\infty$, conditional on k_{t+1}, and then on k_{t+1}. That is, we would choose savings at t, and later the rest. Let us denote by $V(k_t)$ the value of the optimal program from period t for an initial condition k_t:

$$V(k_t) \equiv \max_{\{k_{s+1}\}_{s=t}^{\infty}} \sum_{s=t}^{\infty} \beta^{s-t} F(k_s, k_{s+1}), \text{ s.t. } k_{s+1} \in \Gamma(k_s) \forall s \geq t,$$

where $\Gamma(k_t)$ represents the feasible choice set for k_{t+1} given k_t.[5] That is, V is an indirect utility function, with k_t representing the parameter governing the choices and resulting utility. Then using the maximization-by-steps idea, we can write

$$V(k_t) = \max_{k_{t+1}\in\Gamma(k_t)} \{F(k_t, k_{t+1}) + \max_{\{k_{s+1}\}_{s=t+1}^{\infty}} \sum_{s=t+1}^{\infty} \beta^{s-t} F(k_s, k_{s+1}) \text{ (s.t. } k_{s+1} \in \Gamma(k_s) \forall s \geq t)\}$$

which in turn can be rewritten as

$$\max_{k_{t+1}\in\Gamma(k_t)} \{F(k_t, k_{t+1}) + \beta \max_{\{k_{s+1}\}_{s=t+1}^{\infty}} \{ \sum_{s=t+1}^{\infty} \beta^{s-(t+1)} F(k_s, k_{s+1}) \text{ (s.t. } k_{s+1} \in \Gamma(k_s) \forall s \geq t+1)\}\}.$$

But by definition of V this equals

$$\max_{k_{t+1}\in\Gamma(k_t)} \{F(k_t, k_{t+1}) + \beta V(k_{t+1})\}.$$

So we have:

$$V(k_t) = \max_{k_{t+1}\in\Gamma(k_t)} \{F(k_t, k_{t+1}) + \beta V(k_{t+1})\}.$$

This is the dynamic programming formulation. The derivation was completed for a given value of k_t on the left-hand side of the equation. On the right-hand side, however, we need to know V evaluated at any value for k_{t+1} in order to be able to perform the maximization. If, in other words, we find a V that, using k to denote current capital and k' next period's capital, satisfies

$$V(k) = \max_{k'\in\Gamma(k)} \{F(k, k') + \beta V(k')\}$$

for any value of k, then all the maximizations on the right-hand side are well-defined. This equation is called the Bellman equation, and it is a *functional equation*: the unknown is a function. We use the function g alluded to above to denote the arg max in the functional equation:

$$g(k) = \arg \max_{k'\in\Gamma(k)} \{F(k, k') + \beta V(k')\}$$

or the decision rule for k': $k' = g(k)$. This notation presumes that a maximum exists and is unique; otherwise, g would not be a well-defined function.

This is "close" to a formal derivation of the equivalence between the sequential formulation of the dynamic optimization and its recursive, Bellman formulation. What remains to be done mathematically is to make sure that all the operations above are well-defined. Mathematically, one would want to establish:

[5]The one-sector growth model example would mean that $F(x, y) = u(f(x)-y)$ and that $\Gamma(x) = [0, f(x)]$ (the latter restricting consumption to be non-negative and capital to be non-negative).

- If a function represents the value of solving the sequential problem (for any initial condition), then this function solves the dynamic programming equation (DPE).
- If a function solves the DPE, then it gives the value of the optimal program in the sequential formulation.
- If a sequence solves the sequential program, it can be expressed as a decision rule that solves the maximization problem associated with the DPE.
- If we have a decision rule for a DPE, it generates sequences that solve the sequential problem.

These four facts can be proved, under appropriate assumptions.[6] We omit discussion of details here.

One issue is useful to touch on before proceeding to the practical implementation of dynamic programming: since the maximization that needs to be done in the DPE is finite-dimensional, ordinary Kuhn-Tucker methods can be used, without reference to extra conditions, such as the transversality condition. How come we do not need a transversality condition here? The answer is subtle and mathematical in nature. In the statements and proofs of equivalence between the sequential and the recursive methods, it is necessary to impose conditions on the function V: not any function is allowed. Uniqueness of solutions to the DPE, for example, only follow by resticting V to lie in a restricted space of functions. This or other, related, restrictions play the role of ensuring that the transversality condition is met.

We will make use of some important results regarding dynamic programming. They are summarized in the following:

Facts: Suppose that F is continuously differentiable in its two arguments, that it is strictly increasing in its first argument (and decreasing in the second), strictly concave, and bounded. Suppose that Γ is a nonempty, compact-valued, monotone, and continuous correspondence with a convex graph. Finally, suppose that $\beta \in (0, 1)$. Then

1. There exists a function $V(\cdot)$ that solves the Bellman equation. This solution is unique.
2. It is possible to find V by the following iterative process:
   i. Pick any initial V_0 function, for example $V_0(k) = 0\ \forall k$.
   ii. Find V_{n+1}, for any value of k, by evaluating the right-hand side of (FE) using V_n.

   The outcome of this process is a sequence of functions $\{V_j\}_{j=0}^{\infty}$ which converges to V.
3. V is strictly concave.
4. V is strictly increasing.
5. V differentiable.
6. Optimal behavior can be characterized by a function g, with $k' = g(k)$, that is increasing so long as F_2 is increasing in k.

The proof of the existence and uniqueness part follow by showing that the functional equation's right-hand side is a contraction mapping, and using the contraction mapping theorem. The algorithm for finding V also uses the contraction property. The assumptions needed for these characterizations do not rely on properties of F other than its continuity and boundedness. That is, these results are quite general.

In order to prove that V is increasing, it is necessary to assume that F is increasing and that Γ is monotone. In order to show that V is (strictly) concave it is necessary to assume that F is (strictly) concave and that Γ has a convex graph. Both these results use the iterative algorithm. They essentially require showing that, if the initial guess on V, V_0, satisfies the required property (such as being increasing), then so is any subsequent V_n. These proofs are straightforward.

Differentiability of V requires F to be continuously differentiable and concave, and the proof is somewhat more involved. Finally, optimal policy is a function when F is strictly concave and Γ is convex-valued; under

[6]See Stokey and Lucas (1989).

these assumption, it is also easy to show, using the first-order condition in the maximization, that g is increasing. This condition reads

$$-F_2(k,k') = \beta V'(k').$$

The left-hand side of this equality is clearly increasing in k', since F is strictly concave, and the right-hand side is strictly decreasing in k', since V is strictly concave under the stated assumptions. Furthermore, since the right-hand side is independent of k but the left-hand side is decreasing in k, the optimal choice of k' is increasing in k.

The proofs of all these results can be found in Stokey and Lucas with Prescott (1989).

3.2.1 Connections with finite-horizon problems

Consider the finite-horizon problem

$$\max_{\{c_t\}_{t=0}^{T}} \sum_{t=0}^{T} u(c_t)$$
$$s.t. k_{t+1} + c_t = F(k_t).$$

Although we discussed how to solve this problem in the previous sections, dynamic programming offers us a new solution method. Let $V_n(k)$ denote the present value utility derived from having a current capital stock of k and behaving optimally, if there are n periods left until the end of the world. Then we can solve the problem recursively, or by backwards induction, as follows. If there are no periods left, that is, if we are at $t = T$, then the present value of utility next period will be 0 no matter how much capital is chosen to be saved: $V_0(k) = 0 \; \forall k$. Then once he reaches $t = T - 1$ the consumer will face the following problem:

$$V_1(k) = \max_{k'} \{u[f(k) - k'] + \beta V_0(k')\}.$$

Since $V_0(k') = 0$, this reduces to $V_1(k) = \max_{k'} \{u[f(k) - k']\}$. The solution is clearly $k' = 0$ (note that this is consistent with the result $k_{T+1} = 0$ that showed up in finite horizon problems when the formulation was sequential). As a result, the update is $V_1(k) = u[f(k)]$. We can iterate in the same fashion T times, all the way to V_0, by successively plugging in the updates V_n. This will yield the solution to our problem.

In this solution of the finite-horizon problem, we have obtained an interpretation of the iterative solution method for the infinite-horizon problem: the iterative solution is like solving a finite-horizon problem backwards, for an increasing time horizon. The statement that the limit function converges thus says that the value function of the infinite-horizon problem is the limit of the time-zero value functions of the finite-horizon problems, as the horizon increases to infinity. This also means that the behavior at time zero in a finite-horizon problem becomes increasingly similar to infinite-horizon behavior as the horizon increases.

Finally, notice that we used dynamic programming to describe how to solve a non-stationary problem. This may be confusing, as we stated early on that dynamic programming builds on stationarity. However, if time is viewed as a state variable, as we actually did view it now, the problem can be viewed as stationary. That is, if we increase the state variable from not just including k, but t as well (or the number of periods left), then dynamic programming can again be used.

3.2.2 Solving a parametric dynamic programming problem

Consider the following functional equation:

$$V(k) = \max_{c,\, k'} \{\log c + \beta V(k')\}$$
$$s.t. \; c = Ak^{\alpha} - k'.$$

The budget constraint is written as an equality constraint because we know that preferences represented by the logarithmic utility function exhibit strict monotonicity—goods are always valuable, so they will not be thrown away by an optimizing decision maker. The production technology is represented by a Cobb-Douglass function, and there is full depreciation of the capital stock in every period:

$$\underbrace{F(k,1)}_{Ak^{\alpha}1^{1-\alpha}} + \underbrace{(1-\delta)}_{0}k.$$

A more compact expression can be derived by substitutions into the Bellman equation:

$$V(k) = \max_{k' \geq 0} \{\log[Ak^\alpha - k'] + \beta V(k')\}.$$

We will solve the problem by iterating on the value function. The procedure will be similar to that of solving a T-problem backwards. We begin with an initial "guess" $V_0(k) = 0$, that is, a function that is zero-valued everywhere.

$$\begin{aligned} V_1(k) &= \max_{k' \geq 0} \{\log[Ak^\alpha - k'] + \beta V_0(k')\} \\ &= \max_{k' \geq 0} \{\log[Ak^\alpha - k'] + \beta \cdot 0\} \\ & \max_{k' \geq 0} \{\log[Ak^\alpha - k']\} \end{aligned}$$

This is maximized by taking $k' = 0$. Then

$$V_1(k) = \log A + \alpha \log k.$$

Going to the next step in the iteration,

$$\begin{aligned} V_2(k) &= \max_{k' \geq 0} \{\log[Ak^\alpha - k'] + \beta V_1(k')\} \\ &= \max_{k' \geq 0} \{\log[Ak^\alpha - k'] + \beta[\log A + \alpha \log k']\}. \end{aligned}$$

The first-order condition now reads

$$\frac{1}{Ak^\alpha - k'} = \frac{\beta\alpha}{k'} \Rightarrow k' = \frac{\alpha\beta Ak^\alpha}{1 + \alpha\beta}.$$

We can interpret the resulting expression for k' as the rule that determines how much it would be optimal to save if we were at period $T - 1$ in the finite horizon model. Substitution implies

$$\begin{aligned} V_2(k) &= \log\left[Ak^\alpha - \frac{\alpha\beta Ak^\alpha}{1 + \alpha\beta}\right] + \beta\left[\log A + \alpha \log \frac{\alpha\beta Ak^\alpha}{1 + \alpha\beta}\right] \\ &= \left(\alpha + \alpha^2\beta\right)\log k + \log\left(A - \frac{\alpha\beta A}{1 + \alpha\beta}\right) + \beta \log A + \alpha\beta \log \frac{\alpha\beta A}{1 + \alpha\beta}. \end{aligned}$$

We could now use $V_2(k)$ again in the algorithm to obtain a $V_3(k)$, and so on. We know by the characterizations above that this procedure would make the sequence of value functions converge to some $V^*(k)$. However, there is a more direct approach, using a pattern that appeared already in our iteration.

Let

$$a \equiv \log\left(A - \frac{\alpha\beta A}{1 + \alpha\beta}\right) + \beta \log A + \alpha\beta \log \frac{\alpha\beta A}{1 + \alpha\beta}$$

and

$$b \equiv \left(\alpha + \alpha^2\beta\right).$$

Then $V_2(k) = a + b\log k$. Recall that $V_1(k) = \log A + \alpha \log k$. That is, in the second step what we did was plug in a function $V_1(k) = a_1 + b_1 \log k$, and out came a function $V_2(k) = a_2 + b_2 \log k$. Clearly this implies that if we continue using our iterative procedure, the outcomes $V_3(k)$, $V_4(k)$, ..., $V_n(k)$, will be of the form $V_n(k) = a_n + b_n \log k$ for all n. Therefore, we may already guess that the function to which this sequence is converging has to be of the form:

$$V(k) = a + b\log k.$$

So let's guess that the value function solving the Bellman has this form, and determine the corresponding parameters a, b :

$$V(k) = a + b\log k = \max_{k' \geq 0} \{\log(Ak^\alpha - k') + \beta(a + b\log k')\} \;\; \forall k.$$

Our task is to find the values of a and b such that this equality holds for all possible values of k. If we obtain these values, the functional equation will be solved.

The first-order condition reads:

$$\frac{1}{Ak^{\alpha} - k'} = \frac{\beta b}{k'} \;\Rightarrow k' = \frac{\beta b}{1 + \beta b} Ak^{\alpha}.$$

We can interpret $\dfrac{\beta b}{1 + \beta b}$ as a savings rate. Therefore, in this setup the optimal policy will be to save a constant fraction out of each period's income.

Define

$$LHS \equiv a + b \log k$$

and

$$RHS \equiv \max_{k' \geq 0} \left\{ \log \left(Ak^{\alpha} - k' \right) + \beta \left(a + b \log k' \right) \right\}.$$

Plugging the expression for k' into the RHS, we obtain:

$$\begin{aligned} RHS &= \log \left(Ak^{\alpha} - \frac{\beta b}{1 + \beta b} Ak^{\alpha} \right) + a\beta + b\beta \cdot \log \left(\frac{\beta b}{1 + \beta b} Ak^{\alpha} \right) \\ &= \log \left[\left(1 - \frac{\beta b}{1 + \beta b} \right) Ak^{\alpha} \right] + a\beta + b\beta \cdot \log \left(\frac{\beta b}{1 + \beta b} Ak^{\alpha} \right) \\ &= (1 + b\beta) \log A + \log \left(\frac{1}{1 + b\beta} \right) + a\beta + b\beta \cdot \log \left(\frac{\beta b}{1 + \beta b} \right) + (\alpha + \alpha\beta) \log k. \end{aligned}$$

Setting LHS=RHS, we produce

$$\left\{ \begin{array}{l} a = (1 + b\beta) \log A + \log \left(\dfrac{1}{1 + b\beta} \right) + a\beta + b\beta \cdot \log \left(\dfrac{\beta b}{1 + \beta b} \right) \\ \\ b = (\alpha + \alpha\beta), \end{array} \right.$$

which amounts to two equations in two unknowns. The solutions will be

$$\boxed{b = \frac{\alpha}{1 - \alpha\beta}}$$

and, using this finding,

$$a = \frac{1}{1 - \beta} \cdot \left[(1 + b\beta) \log A + b\beta \log (b\beta) - (1 + \beta) \log (1 + b\beta) \right]$$

so that

$$\boxed{a = \frac{1}{1 - \beta} \cdot \frac{1}{1 - \alpha\beta} \cdot \left[\log A + \log (1 - \alpha\beta) + \alpha\beta \cdot \log (\alpha\beta) \right]}$$

Going back to the savings decision rule, we have:

$$\begin{aligned} k' &= \frac{b\beta}{1 - b\beta} Ak^{\alpha} \\ k' &= \alpha\beta Ak^{\alpha} \end{aligned}$$

If we let y denote income, that is, $y \equiv Ak^{\alpha}$, then $k' = \alpha\beta \cdot y$. This means that the optimal solution to the path for consumption and capital is to save a constant fraction $\alpha\beta$ of income.

This setting, we have now shown, provides a microeconomic justification to a constant savings rate, like the one assumed by Solow. It is a very special setup however, one that is quite restrictive in terms of functional forms. Solow's assumption cannot be shown to hold generally.

We can also visualize the dynamic behavior of capital:

Figure 3: The decision rule in our parameterized model

3.2.3 A more complex example

We will now look at a slightly different growth model and try to put it in recursive terms. Our new problem is:

$$\max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \cdot u(c_t)$$
$$s.t. \quad c_t + i_t = F(k_t)$$

subject to the assumption is that capital depreciates fully in <u>two</u> periods, and doesn't depreciate at all before that. Then the law of motion for capital, given a sequence of investment $\{i_t\}_{t=0}^{\infty}$ is given by:

$$k_t = i_{t-1} + i_{t-2}.$$

Then $k_0 = i_{-1} + i_{-2}$: there are two initial conditions i_{-1} and i_{-2}.

The recursive formulation for this problem is:

$$\begin{aligned} V(i_{-1},\ i_{-2}) = \quad & \max_{c,\ i} \{u(c) + V(i,\ i_{-1})\} \\ & s.t. \\ c = f(i_{-1} + i_{-2}) - i. \end{aligned} \tag{FE}$$

Notice that there are two state variables in this problem. That is unavoidable here; there is no way of summarizing what one needs to know at a point in time with only one variable. For example, the total capital stock in the current period is not informative enough, because in order to know the capital stock next period we need to know how much of the current stock will disappear between this period and the next. Both i_{-1} and i_{-2} are natural state variables: they are predetermined, they affect outcomes and utility, and neither is redundant: the information they contain cannot be summarized in a simpler way.

3.2.4 The functional Euler equation

In the sequentially formulated maximization problem, the Euler equation turned out to be a crucial part of characterizing the solution. With the recursive strategy, an Euler equation can be derived as well. Consider again

$$V(k) = \max_{k' \in \Gamma(k)} \{F(k,\ k') + \beta V(k')\}.$$

As already pointed out, under suitable assumptions, this problem will result in a function $k' = g(k)$ that we call decision rule, or policy function. By definition, then, we have

$$V(k) = F(k,\ g(k)) + \beta V[g(k)].$$

Moreover, $g(k)$ satisfies the first-order condition

$$F_2(k,\ k') + \beta V'(k') = 0, \tag{FOC}$$

assuming an interior solution. Evaluating at the optimum, i.e., at $k' = g(k)$, we have

$$F_2(k,\ g(k)) + \beta V'[g(k)] = 0.$$

This equation governs the intertemporal tradeoff. One problem in our characterization is that $V'(\cdot)$ is not known: in the recursive strategy, it is part of what we are searching for. However, although it is not possible in general to write $V(\cdot)$ in terms of primitives, one can find its derivative. Using the equation for $V(k)$ above, one can differentiate both sides with respect to k, since the equation holds for all k and, again under some assumptions stated earlier, is differentiable. We obtain

$$V'(k) = F_1 [k,\ g(k)] + \underbrace{g'(k) \cdot \{F_2 [k,\ g(k)] + \beta V' [g(k)]\}}_{\text{indirect effect through optimal choice of } k'} .$$

From the first-order condition, this reduces to

$$V'(k) = F_1 [k,\ g(k)] ,$$

which again holds for all values of k. The indirect effect thus disappears: this is an application of a general result known as the envelope theorem.

Updating, we know that $V' [g(k)] = F_1 [g(k),\ g (g[k])]$ also has to hold. The first order condition can now be rewritten as follows:

$$F_2 [k,\ g(k)] + \beta F_1 [g(k),\ g (g(k))] = 0 \ \forall k.$$

This is the Euler equation stated as a functional equation: it does not contain the unknowns k_t, k_{t+1}, and k_{t+2}; recall our previous Euler equation formulation

$$F_2 [k_t,\ k_{t+1}] + \beta F_1 [k_{t+1},\ k_{t+2}] = 0.$$

Now instead, the unknown is the function g. That is, under the recursive formulation, the Euler Equation has turned into a functional equation.

The previous discussion suggests that a third way of searching for a solution to the dynamic problem is to consider the functional Euler equation, and solve it for the function g. We have previously seen that we can (i) look for sequences solving a nonlinear difference equation plus a transversality condition; and (ii) we can solve a Bellman (functional) equation for a value function. The functional Euler equation approach is, in some sense, somewhere in between the two previous approaches. It is based on an equation expressing an intertemporal tradeoff, but it applies more structure than our previous Euler equation. There, a transversality condition needed to be invoked in order to find a solution. Here, we can see that the recursive approach provides some extra structure: it tells us that the optimal sequence of capital stocks needs to be connected using a stationary function. One problem is that the functional Euler equation does not in general have a unique solution for g. It might, for example, have two solutions. This multiplicity is less severe, however, than the multiplicity in a second-order difference equation without a transversality condition: there, there are infinitely many solutions.

The functional Euler equation approach is often used in practice in solving dynamic problems numerically. We will return to this equation below.

4 Steady states and dynamics under optimal growth

We will now study, in more detail, the model where there is only one type of good, that is, only one production sector: the one-sector optimal growth model. This means that we will revisit the Solow model under the assumption that savings are chosen optimally. Will, as in Solow's model, output and all other variables converge to a steady state? It turns out that the one-sector optimal growth model does produce global convergence under fairly general conditions. This can be proven analytically. If the number of sectors increases, however, global convergence may not occur. However, in practical applications, where the parameters describing different sectors are chosen so as to match data, it has proven difficult to find examples where global convergence does not apply.

We thus consider preferences of the type

$$\sum_{t=0}^{\infty} \beta^t \cdot u (c_t)$$

and production given by

$$c_t + k_{t+1} = f(k_t)$$

Figure 4: The determination of steady state

where

$$f(k_t) = F(k_t,\ N) + (1-\delta)\,k_t$$

for some choice of N and δ (which are exogenous in the setup we are looking at). Under standard assumptions (namely strict concavity, $\beta < 1$, and conditions ensuring interior solutions), we obtain the Euler equation:

$$u'(c_t) = \beta \cdot u'(c_{t+1}) \cdot f'(k_{t+1})\,. \tag{EE}$$

A *steady state* is a "constant solution":

$$\begin{aligned} k_t &= k^* \forall t \\ c_t &= c^* \forall t. \end{aligned}$$

This constant sequence $\{c_t\}_{t=0}^{\infty} = \{c^*\}_{t=0}^{\infty}$ will have to satisfy:

$$u'(c^*) = \beta \cdot u'(c^*) \cdot f'(k^*)\,.$$

Here $u'(c^*) > 0$ is assumed, so this reduces to

$$\beta \cdot f'(k^*) = 1.$$

This is the key condition for a steady state in the one-sector growth model. It requires that the gross marginal productivity of (return to) capital equal the gross discount rate $(1/\beta)$.

Suppose $k_0 = k^*$. We first have to ask whether $k_t = k^*\ \forall t$—a solution to the steady-state equation—will solve the maximization problem. The answer is clearly yes, provided that both the first order *and* the transversality conditions are met! The first order conditions are met by construction, with consumption defined by

$$c^* = f(k^*) - k^*.$$

The transversality condition requires

$$\lim_{t\to\infty} \beta^t \cdot F_1[k_t,\ k_{t+1}] \cdot k_t = 0. \tag{TVC}$$

Evaluated at the proposed sequence, this condition becomes

$$\lim_{t\to\infty} \beta^t \cdot F_1[k^*,\ k^*] \cdot k^* = 0,$$

and since $F_1[k^*,\ k^*] \cdot k^*$ is a finite number, with $\beta < 1$, the limit clearly is zero and the condition is met. Therefore we can conclude that the stationary solution $k_t = k^*\ \forall t$ does maximize the objective function. If f is strictly concave, then k^* is the unique strictly positive solution. It remains to verify that there is indeed one solution. We will get back to this in a moment.

Graphically, concavity of $f(k)$ implies that $\beta \cdot f'(k)$ will be a positive, decreasing function of k, and it will intersect 1 at one point:

4.1 Global convergence

In order to derive precise results for global convergence, we now make the following assumptions on primitives:

(i) u and f are strictly increasing, strictly concave, and continuously differentiable.

(ii) $f(0) = 0$, $\lim_{k \to 0} f'(k) = \infty$, and $\lim_{k \to \infty} f'(k) \equiv b < 1$.[7]

(iii) $\lim_{c \to 0} u'(c) = \infty$.

(iv) $\beta \in (0,\ 1)$.

We have the problem:

$$V(k) = \max_{k' \in [0, f(k) - k']} \{u[f(k) - k'] + \beta V(k')\}$$

leading to $k' = g(k)$ satisfying the first-order condition

$$u'[f(k) - k'] = \beta V'(k').$$

Notice that we are assuming an interior solution. But this assumption is valid, since assumptions (ii) and (iii) guarantee interiority.

In order to prove global convergence to a unique steady state, we will first state (and rederive) some properties of the policy function g and of steady states. We will then use these properties to rule out anything but global convergence.

Properties of g:

Property 1 $g(k)$ is single-valued for all k.

This follows from strict concavity of u and V (recall the theorem we stated previously).

Property 2 $g(k)$ is strictly increasing.

We argued this informally in the previous section. The argument is as follows. Consider the first-order condition:

$$u'[f(k) - k'] = \beta V'(k').$$

$V'(\cdot)$ is decreasing, since $V(\cdot)$ is strictly concave due to the assumptions on u and f. Define

$$\begin{aligned} RHS(k, k') &= u'[f(k) - k'] \\ LHS(k') &= \beta V'(k'). \end{aligned}$$

Let $\tilde{k} > k$. Then $f(\tilde{k}) - k' > f(k) - k'$. Strict concavity of u implies that $u'\left[f(\tilde{k}) - k'\right] < u'[f(k) - k']$. Hence we have that

$$\tilde{k} > k \Rightarrow RHS(\tilde{k}, k') < RHS(k, k').$$

As a consequence, the $LHS(k')$ must decrease to satisfy the first order condition. Since $V'(\cdot)$ is decreasing, this will happen only if k' increases. This shows that $\tilde{k} > k \Rightarrow g(\tilde{k}) > g(k)$.

This can also be viewed as an application of the implicit function theorem. Define

$$H(k, k') \equiv u'[f(k) - k'] - \beta V'(k') = 0.$$

[7] It is not necessary for the pursuing arguments to assume that $\lim_{k \to 0} f'(k) = \infty$; it would work if the limit is strictly greater than 1.

Then

$$
\begin{aligned}
\frac{\partial k'}{\partial k} &= -\frac{\dfrac{\partial H(k,k')}{\partial k}}{\dfrac{\partial H(k,k')}{\partial k'}} \\
\frac{\partial k'}{\partial k} &= -\frac{u''\left[f(k)-k'\right]}{-u''\left[f(k)-k'\right]-\beta V''(k')} \\
\frac{\partial k'}{\partial k} &= \frac{u''\left[f(k)-k'\right]}{u''\left[f(k)-k'\right]+\beta V''(k')} > 0,
\end{aligned}
$$

where the sign follows from the fact that since u and V are strictly concave, both the numerator and the denominator of this expression have negative signs. This derivation is heuristic; we have assumed here that V is twice differentiable. It turns out that there is a theorem telling us that (under some side conditions that we will not state here) V will indeed be twice differentiable, given that u and f are both twice differentiable, but it is beyond the scope of the present analysis to discuss this theorem in more detail.

The economic intuition behind g being increasing is simple. There is an underlying presumption of normal goods behind our assumptions: strict concavity and additivity of the different consumption goods (over time) amounts to assuming that the different goods are normal goods. Specifically, consumption in the future is a normal good. Therefore, a larger initial wealth commands larger savings.

Property 3 $g(k)$ is continuous.

This property follows from the Theorem of the Maximum under concavity, which you are invited to study in more detail in your mathematics text!

Property 4 $g(0) = 0$.

This follows from $f(k) - k' \geq 0$ and $f(0) = 0$.

Property 5 There exists a $\bar{k}$ such that $g(k) \leq \bar{k}$ for all $k < \bar{k}$. Moreover, $\bar{k}$ exceeds $(f')^{-1}(1/\beta)$.

The first part follows from feasibility: because consumption cannot be negative, k' cannot exceed $f(k)$. Our assumptions on f then guarantee that $f(k) < k$ for high enough values of k: the slope of f approaches a number less than 1 as k goes to infinity. So $g(k) < k$ follows. The characterization of $\bar{k}$ follows from noting (i) that $\bar{k}$ must be above the value that maximizes $f(k) - k$, since $f(k)$ is above k for very small values of k and f is strictly concave and (ii) that therefore $\bar{k} > (f')^{-1}(1) > (f')^{-1}(1/\beta)$.

We saw before that a necessary and sufficient condition for a steady state is that $\beta \cdot f'(k^*) = 1$. Now let us consider different possibilities for the decision rule. The following figure shows three decision rules which are all increasing and continuous.

All candidate decision rules start at the origin, as required by the assumption that $f(0) = 0$ and that investment is nonnegative. They all stay in the bounded set $[0, \bar{k}]$. Line 1 has three different solutions to the steady-state condition $k' = k$. We do not graph the possibility of more than three (or two) steady states; these will be ruled out with the same argument as line 1. Line 2 has only one steady state. Line 3 has no positive steady state. There are no other possibilities.

We need to investigate further the number of steady states that our model will allow.

Figure 5: Different decision rule candidates

Figure 6: How many steady states obtain?

Figure **??** shows that a single steady state must obtain, due to strict concavity of u; in fact, this result was derived above. This rules out line 1 in Figure **??**, and it also rules out any decision rule with more than one positive crossing.

Line 3, with no positive steady state, can be ruled out by noting that the steady state requirement $\beta \cdot f'(k^*) = 1$, together with Property 5, mean that there will be a strictly positive steady state in the interval $(0, \bar{k})$. Therefore the only possibility is:

We can see that the blue line is above the 45^odegree line to the left of k^*, and below to the right. This implies that the model dynamics exhibit global convergence.

The convergence will not occur in finite time. For it to occur in finite time, the decision rule would have to be flat through the steady state point. However, we have established that it is strictly increasing (Property 2). A case with a flat decision rule is depicted in the following graph; the decision rule in that graph is not only impossible because it has a flat part; it also has its stationary point at $\bar{k}$, which we know is not possible. The graph would imply zero consumption for the entire right (flat) region.

4.2 Dynamics: the speed of convergence

What can we say about the time it takes to reach the steady state? The speed of global convergence will depend on the shape of $g(k)$, as the following chart shows:

Capital will approach the steady state level more rapidly (that is, in "a smaller number of steps") along trajectory number 2: it will have a faster *speed of convergence*. There is no simple way to summarize, in a quantitative way, the speed of convergence for a general decision rule. However, for a limited class of decision rules—the linear (or affine) rules—it can be measured simply by looking at the slope. This is an important case, for it can be used locally to approximate the speed of convergence around the steady state k^*. The argument for this is simple: the accumulation path will spend infinite time arbitrarily close to the steady state, and in a very small region a continuous function can be arbitrarily well approximated by a linear function, using the first-order Taylor expansion of the function. That is, for any capital accumulation path, we will be able to approximate the speed of convergence arbitrarily well as time passes. If the starting point is far from the steady state, we will make mistakes that might be large initially, but these mistakes will

Figure 7: The remaining candidate

Figure 8: Satiation

Figure 9: Different speeds of convergence

become smaller and smaller and eventually become unimportant. Moreover, if one uses parameter values that are, in some sense, realistic, it turns out that the resulting decision rule will be quite close to linear.

In this section, we will state a general theorem with properties for dynamic systems of a general size: that is, we will be much more general than the one-sector growth model. In fact, it is possible with the methods we describe here to obtain the key information about local dynamics for any dynamic system. The global convergence theorem, in contrast, applies only for the one-sector growth model.

The first-order Taylor series expansion of the decision rule gives

$$\begin{aligned} k' &= g(k) \approx \underbrace{g(k^*)}_{k^*} + g'(k^*) \cdot (k - k^*) \\ \underbrace{k' - k^*}_{\text{Next period's gap}} &= g'(k^*) \cdot \underbrace{(k - k^*)}_{\text{Current gap}} . \end{aligned}$$

This shows that we may interpret $g'(k^*)$ as a measure of the rate of convergence (or rather, its inverse). If $g'(k^*)$ is very close to zero, convergence is fast—the gap is decreasing significantly each period.

4.2.1 Linearization for a general dynamic system

The task is now to find $g'(k^*)$ by *linearization*. We will use the Euler equation and linearize it. This will lead to a difference equation in k_t. One of the solutions to this difference equation will be the one we are looking for. Two natural questions arise: 1) How many convergent solutions are there around k^*? 2) For the convergent solutions, is it valid to analyze a linear difference equation as a proxy to their convergence speed properties? The first of these questions is the key to the general characterization of dynamics. The second question is a mathematical one, making the sense of approximation precise.

Both questions are addressed by the following theorem, which applies to a general dynamic system (that is, not only those coming from economic models):

Theorem 3 *Let $x_t \in \Re^n$. Given $x_{t+1} = h(x_t)$ with a stationary point $\bar{x} : \bar{x} = h(\bar{x})$. If*

1. *h is continuously differentiable with Jacobian $H(\bar{x})$ around $\bar{x}$.*
2. *$I - H(\bar{x})$ is non-singular.*

then there is a set of initial conditions x_0, of dimension equal to the number of eigenvalues of $H(\bar{x})$ that are less than 1 in absolute value, for which $x_t \to \bar{x}$.

We will describe how to use this theorem with a few examples.

Example 4 **($n = 1$)** *There is only one eigenvalue: $\lambda = h'(\bar{x})$*

1. *$|\lambda| \geq 1 \Rightarrow$ no initial condition leads to x_t converging to $\bar{x}$.*

   In this case, only for $x_0 = \bar{x}$ will the system stay in $\bar{x}$.

2. $|\lambda| < 1 \Rightarrow x_t \to \bar{x}$ *for any value of* x_0.

Example 5 **($n = 2$)** *1.* $|\lambda_1|, |\lambda_2| \geq 1$ $\Rightarrow$ *No initial condition* x_0 *leads to convergence.*

2. $|\lambda_1| < 1$, $|\lambda_2| \geq 1 \Rightarrow$ *Dimension of* x_0*'s leading to convergence is* 1. *This is called "saddle path stability".*

3. $|\lambda_1|, |\lambda_2| < 1$ $\Rightarrow$ *Dimension of* x_0*'s leading to convergence is* 2. $x_t \to \bar{x}$ *for any value of* x_0.

The examples, hopefully, describe how a general dynamic system behaves. It does not yet, however, quite settle the issue of convergence. In particular, the set of initial conditions leading to convergence must be given an economic meaning. Is any initial condition possible in a given economic model? Typically no: for example, the initial capital stock in an economy may be given, and thus we have to restrict the set of initial conditions to those respecting the initial capital stock.

We will show below that an economic model has dynamics that can be reduced to a vector difference equation of the form of the one described in the above theorem. In this description, the vector will have a subset of true state variables (like capital); the remainder of the vector consists of various control, or other, variables that are there in order that the system can be put in first-order form.

More formally, let the number of eigenvalues less than 1 in absolute value be denoted by m. This is the dimension of the set of initial x_0's leading to $\bar{x}$. We may interpret m as the *degrees of freedom*. Let the number of economic restrictions on initial conditions be denoted by $\hat{m}$. These are the restrictions emanating from physical (and perhaps other) conditions in our economic model. Notice that an interpretation of this is that we have $\hat{m}$ equations and m unknowns. Then the issue of convergence boils down to the following cases.

1. $m = \hat{m} \Rightarrow$ there is a *unique* convergent solution to the difference equation system.
2. $m < \hat{m} \Rightarrow$ No convergent solution obtains.
3. $m > \hat{m} \Rightarrow$ There is "indeterminacy", that is, "many" solutions (how many? dim $= \hat{m} - m$).

4.2.2 The economic model

We now describe in detail how the linearization procedure works. The example comes from the one-sector growth model, but the general outline is the same for all economic models.

1. Derivation of EE: $F(k_t, k_{t+1}, k_{t+2}) = 0$

$$u'[f(k_t) - k_{t+1}] - \beta \cdot u'[f(k_{t+1}) - k_{t+2}] \cdot f'(k_{t+1}) = 0.$$

k^* is a steady state $\Leftrightarrow F(k^*, k^*, k^*) = 0$.

2. Linearize the EE: Define $\hat{k}_t = k_t - k^*$ and derive

$$a_2 \cdot \hat{k}_{t+2} + a_1 \cdot \hat{k}_{t+1} + a_0 \cdot \hat{k}_t = 0.$$

3. Write the EE as a first-order system: A difference equation of any order can be written as a first order difference equation by using vector notation: Define $\mathrm{x}_t = \begin{pmatrix} \hat{k}_{t+1} \\ \hat{k}_t \end{pmatrix}$

$$\mathrm{x}_{t+1} = H \cdot \mathrm{x}_t. \tag{6}$$

4. Find the solution to the first-order system:

$$x_t = c_1 \cdot \lambda_1^t \cdot \mathrm{v}_1 + c_2 \cdot \lambda_2^t \cdot \mathrm{v}_2,$$

where c_1 and c_2 are constants to be determined, λ_1 and λ_2 are (distinct) eigenvalues of H, and v_1 and v_2 are eigenvectors associated with these eigenvalues.

5. Determine the constants: Use the information about state variables and initial conditions to find c_1 and c_2. In this case, x consists of one state variable and one lagged state variable, the latter used only for the reformulation of the dynamic system. Therefore, we have one initial condition for the system, given by k_0; this amounts to one restriction on the two constants. The set of initial conditions for x_0 in our economic model has therefore been reduced to one dimension. Finally, we are looking for convergent solutions. If one of the two eigenvalues is larger than one in absolute value, this means that we need to set the corresponding constant to zero. Consequently, since not only k_0 but also k_1 are now determined (i.e., both elements of x_0), and our system is fully determined: all future values of k (or x) can be obtained.

   If both eigenvalues are larger than one, the dynamics will not have convergence to the steady state: only if the system starts at the steady state will it remain there.

   If both eigenvalues are less than one, we have no way of pinning down the remaining constant, and the set of converging paths will remain of one dimension. Such indeterminacy—effectively an infinite number of solutions to the system—will not occur in our social planning problem, because its being one single maximization problem will (under strict concavity) guarantee that the set of solutions is a singleton. However, in equilibrium systems that are not derived from a planning problem (perhaps because the equilibrium is not Pareto optimal, as we shall see below), it is possible with indeterminacy.

   The typical outcome in our one-sector growth model is $0 < \lambda_1 < 1$ and $\lambda_2 > 1$. This implies $m = 1$ (saddle path stability); then the convergent solution has $c_2 = 0$. In other words, the economics of our model dictate that the number of restrictions we have on the initial conditions is one, namely the (given) initial level of capital, k_0. That is, $\hat{m} = 1$. Therefore, $m = \hat{m}$, so there is a unique convergent path for each k_0 (close to k^*).

   Then c_1 and c_2 are determined by setting $c_2 = 0$ (so that the path is convergent) and solving for the value of c_1 such that if $t = 0$, then k_t is equal to the given level of initial capital, k_0.

We now implement these steps in detail.

First, we need to solve for H. Let's go back to

$$u'\left[f(k_t) - k_{t+1}\right] - \beta \cdot u'\left[f(k_{t+1}) - k_{t+2}\right] \cdot f'\left(k_{t+1}\right) = 0.$$

In order to linearize, we take derivatives of this expression with respect to k_t, k_{t+1} and k_{t+2}, and evaluate them at k^*. We obtain

$$\beta \cdot u''(c^*) \cdot f'(k^*) \cdot \hat{k}_{t+2} - \left[u''(c^*) + \beta \cdot u''(c^*) \cdot \left[f'(k^*)\right]^2 + \beta \cdot u'(c^*) \cdot f''(k^*)\right] \cdot \hat{k}_{t+1} + u''(c^*) \cdot f'(k^*) \cdot \hat{k}_t = 0.$$

We can use the steady-state fact that $\beta \cdot f'(k^*) = 1$ to simplify this expression:

$$u''(c^*) \cdot \hat{k}_{t+2} - \left[u''(c^*) + \beta^{-1} \cdot u''(c^*) + u'(c^*) \cdot \left[f'(k^*)\right]^{-1} \cdot f''(k^*)\right] \cdot \hat{k}_{t+1} + \beta^{-1} \cdot u''(c^*) \cdot \hat{k}_t = 0.$$

Dividing through by $u''(c^*)$, we arrive at

$$\hat{k}_{t+2} - \left[1 + \frac{1}{\beta} + \frac{u'(c^*)}{u''(c^*)} \cdot \frac{f''(k^*)}{f'(k^*)}\right] \cdot \hat{k}_{t+1} + \frac{1}{\beta} \cdot \hat{k}_t = 0.$$

Then

$$\begin{pmatrix} \hat{k}_{t+2} \\ \hat{k}_{t+1} \end{pmatrix} = H \cdot \begin{pmatrix} \hat{k}_{t+1} \\ \hat{k}_t \end{pmatrix}$$

with

$$H = \begin{pmatrix} 1 + \frac{1}{\beta} + \frac{u'(c^*)}{u''(c^*)} \cdot \frac{f''(k^*)}{f'(k^*)} & -\frac{1}{\beta} \\ 1 & 0 \end{pmatrix}.$$

This is a 2^{nd} order difference equation. Notice that H delivers $\hat{k}_{t+1} = \hat{k}_{t+1}$, so the vector representation of the system is correct. Now we need to look for the eigenvalues of H, from the characteristic polynomial given by

$$|H - \lambda I| = 0.$$

As an interlude before solving for the eigenvalues, let us now motivate the general solution to the linear system above with an explicit derivation from first principles. We can diagonalize H as follows:

$$H = \mathrm{V} \cdot \Lambda \cdot \mathrm{V}^{-1} \;\Rightarrow\; \Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}.$$

Recall that

$$\mathrm{x}_{t+1} = H \cdot \mathrm{x}_t.$$

A change of variables will help us get the solution to this system. First premultiply both sides by V^{-1}:

$$\begin{aligned} \mathrm{V}^{-1} \cdot \mathrm{x}_{t+1} &= \mathrm{V}^{-1} \cdot H \cdot \mathrm{x}_t \\ &= \mathrm{V}^{-1} \cdot \mathrm{V} \cdot \Lambda \cdot \mathrm{V}^{-1} \cdot \mathrm{x}_t \\ &= \Lambda \cdot \mathrm{V}^{-1} \cdot \mathrm{x}_t. \end{aligned}$$

Let $\mathrm{z}_t \equiv \mathrm{V}^{-1} \cdot \mathrm{x}_t$; $\mathrm{z}_{t+1} \equiv \mathrm{V}^{-1} \cdot \mathrm{x}_{t+1}$; then

$$\begin{aligned} \mathrm{z}_{t+1} &= \Lambda \cdot \mathrm{z}_t \\ \mathrm{z}_t &= \Lambda^t \cdot \mathrm{z}_0 \\ \mathrm{z}_{1t} &= c_1 \cdot \lambda_1^t = \mathrm{z}_{10} \cdot \lambda_1^t \\ \mathrm{z}_{2t} &= \mathrm{z}_{20} \cdot \lambda_2^t. \end{aligned}$$

We can go back to x_t by premultiplying z_t by the eigenvector V:

$$\begin{aligned} \mathrm{x}_t &= \mathrm{V} \cdot \mathrm{z}_t \\ &= \mathrm{V} \cdot \begin{pmatrix} \mathrm{z}_{1t} \\ \mathrm{z}_{2t} \end{pmatrix} \\ &= c_1 \cdot \lambda_1^t \cdot \begin{pmatrix} \mathrm{V}_{11} \\ \mathrm{V}_{21} \end{pmatrix} + c_2 \cdot \lambda_2^t \cdot \begin{pmatrix} \mathrm{V}_{12} \\ \mathrm{V}_{22} \end{pmatrix} \\ &= \begin{pmatrix} \hat{k}_{t+1} \\ \hat{k}_t \end{pmatrix}. \end{aligned}$$

The solution, thus, must be of the form

$$\hat{k}_t = \widehat{c}_1 \cdot \lambda_1^t + \widehat{c}_2 \cdot \lambda_2^t,$$

where $\widehat{c}_1 \equiv \dfrac{c_1}{\lambda}$.

To go back to finding the eigenvalues in our specific setting, we thus use $|H - \lambda I| = 0$ to obtain

$$\begin{aligned} \begin{vmatrix} 1 + \dfrac{1}{\beta} + \dfrac{u'}{u''} \cdot \dfrac{f''}{f'} - \lambda & -\dfrac{1}{\beta} \\ 1 & -\lambda \end{vmatrix} &= 0 \\ \lambda^2 - \left[1 + \frac{1}{\beta} + \frac{u'}{u''} \cdot \frac{f''}{f'}\right] \cdot \lambda + \frac{1}{\beta} &= 0, \end{aligned}$$

where u', u'', f', f'' denote the corresponding derivatives evaluated at k^*. Let

$$F(\lambda) \equiv \lambda^2 - \left[1 + \frac{1}{\beta} + \frac{u'}{u''} \cdot \frac{f''}{f'}\right] \cdot \lambda + \frac{1}{\beta}.$$

This is a continuous function of λ, and

$$\begin{aligned} F(0) &= \frac{1}{\beta} > 0 \\ F(1) &= -\frac{u'}{u''} \cdot \frac{f''}{f'} < 0. \end{aligned}$$

Therefore, the mean value theorem implies that $\exists \lambda_1 \in (0,1) : F(\lambda_1) = 0$. That is, one of the eigenvalues is positive and smaller than one. Since $\lim_{\lambda \to \infty} F(\lambda) = +\infty > 0$, the other eigenvalue ($\lambda_2$) must also be positive and larger than 1.

We thus see that a convergent solution to the system requires $c_2 = 0$. The remaining constant, c_1, will be determined from

$$\begin{aligned} \hat{k}_t &= \widehat{c}_1 \cdot \lambda_1^t \\ \hat{k}_0 &\equiv k_0 - k^* \\ &\Rightarrow \widehat{c}_1 = k_0 - k^*. \end{aligned}$$

The solution, therefore, is

$$k_t = k^* + \lambda_1^t \cdot (k_0 - k^*).$$

Recall that

$$k_{t+1} - k^* = g'(k^*) \cdot (k_t - k^*).$$

Analogously, in the linearized system,

$$k_{t+1} - k^* = \lambda_1 \cdot (k_t - k^*).$$

It can thus be seen that the eigenvalue λ_1 has a particular meaning: it measures the (inverse of the) rate of convergence to the steady state.

As a second illustration, suppose we were looking at the larger system

$$\begin{aligned} k_t &= c_1 \cdot \lambda_1^t + c_2 \cdot \lambda_2^t + c_3 \cdot \lambda_3^t + c_4 \cdot \lambda_4^t, \\ & k_0 \text{being given.} \end{aligned}$$

That is, some economic model with a single state variable leads to a third-order difference equation. If only one eigenvalue λ_1 has $|\lambda_1| < 1$, then there is a unique convergent path leading to the steady state. This means that c_2, c_3, c_4, will need to be equal to zero (choosing the subscript 1 to denote the eigenvalue smaller than 1 in absolute value is arbitrary, of course).

In contrast, if there were, for example, two eigenvalues λ_1, λ_2 with $|\lambda_1|, |\lambda_2| < 1$, then we would have $m = 2$ (two "degrees of freedom"). But there is only one economic restriction, namely k_0 given. That is, $\hat{m} = 1 < m$. Then there would be many convergent paths satisfying the sole economic restriction on initial conditions: the system would be indeterminate.

5 Competitive Equilibrium in Dynamic Models

It is now time to leave pure maximization setups where there is a planner making all decisions and move on to market economies. What *economic arrangement*, or what *allocation mechanism*, will be used in the model economy to talk about decentralized, or at least less centralized, behavior? Of course, different physical environments may call for different arrangements. Although many argue that the modern market economy is not well described by well-functioning markets due to the presence of various frictions (incomplete information, externalities, market power, and so on), it still seems a good idea to build the frictionless economy first, and use it as a benchmark from which extensions can be systematically built and evaluated. For a frictionless economy, competititive equilibrium analysis seems suitable.

One issue is what the population structure will be. We will first look at the infinite-horizon (dynastic) setup. The generalization to models with overlapping generations of consumers will come later on. Moreover, we will, whenever we use the competitive equilibrium paradigm, assume that there is a "representative consumer". That is to say that we think of there as being a large (truly infinite, perhaps) number of consumers in the economy who are all identical. Prices of commodities will then have to adjust so that markets clear; this will typically mean (under appropriate strict concavity assumptions) that prices will make all these consumers make the same decisions: prices will have to adjust so that consumers do not interact. For example, the dynamic model without production gives a trivial allocation outcome: the consumer consumes the endowment of every product. The competitive mechanism ensures that this outcome is achieved by prices being set so that the consumer, when viewing prices as beyond his control, chooses to consume no more and no less than his endowments.

For a brief introduction, imagine that the production factors (capital and labor) were owned by many individual *households*, and that the technology to transform those factors into consumption goods was operated by *firms*. Then households' decisions would consist of the amount of factors to provide to firms, and the amount of consumption goods to purchase from them, while firms would have to choose their production volume and factor demand.

The device by which sellers and buyers (of factors and of consumption goods) are driven together is the *market*; clearly this one brings with it the associated concept of *prices*. By equilibrium we mean a situation such that for some given *prices*, individual households' and firms' decisions show an aggregate consistency: the amount of factors that suppliers are willing to supply equals the amount that producers are willing to take, and the same for consumption goods—we say that *markets clear*. The word "competitive" indicates that we are looking at the perfect competition paradigm, as opposed to economies in which firms might have some sort of "market power".

Somewhat more formally, a competitive equilibrium is a vector of prices and quantities that verify certain properties related to the aggregate consistency of individual decisions mentioned above. These properties are:

1. Households choose quantities so as to maximize the level of utility attained given their "wealth" (factor ownership evaluated at the given prices). When making decisions, households take prices as given parameters. The maximum monetary value of goods that households are able to purchase given their wealth is called the budget constraint.

2. The quantity choice is "feasible". By this we mean that the aggregate amount of commodities that individual decision makers have chosen to demand can be produced with the available technology using the amount of factors that suppliers are willing to supply. Notice that this supply is in turn determined by the remuneration to factors—that is, their price. Therefore this second condition is nothing but the requirement that *markets clear*.

3. Firms chose the production volume that maximizes their profits at the given prices.

For dynamic economic setups, we need to specify how trade takes place over time: are the economic agents using assets (and, if so, what kinds of assets)? Often, it will be possible to think of several different economic arrangements for the same physical environment that all give rise to the same final allocations. It will be illustrative to consider, for example, both the case when firms rent their inputs from consumers every period, and thus do not need an intertemporal perspective (and hence assets) to fulfill their profit maximization objective, and the case when they buy and own the long-lived capital they use in production, and hence need to consider the relative values of profits in different periods.

Also, in dynamic competitive equilibrium models, as in the maximization sections above, mathematically there are two alternative procedures: equilibria can be defined and analyzed in terms of (infinite) sequences, and they can be expressed recursively, using functions. We will look at both, starting with the former. For each approach, we will consider both different possible specific arrangements, and we will proceed using examples: we will typically consider an example without production (an "endowment economy") and the neoclassical growth model. Later applied chapters will feature many examples of other setups.

5.1 Sequential competitive equilibrium

The central question is the one of determining the set of commodities that are traded. The most straightforward extension of standard competitive analysis to dynamic models is perhaps the conceptually most abstract one: simply let goods be dated (so that, for example, in a one-good per date context, there is an infinite sequence of commodities: consumption at 0, consumption at 1, and so on) and, like in a static model, let the trade in all these commodities take place once and for all. We will call this setup the date-0 (or Arrow-Debreu-McKenzie) arrangement. In this arrangement, there is no need for assets. If, for example, a consumer needs to consume both in periods 0 and in future periods, the consumer would buy (rights to) future consumption goods at the beginning of time, perhaps in exchange for current labor services, or promises of future labor services. Any activity in the future would then be a mechanical carrying out of all the promises made at time zero.

An alternative setup is one with assets: we will refer to this case as one with sequential trade. In such a case, assets are used by one or more agents, and assets are traded every period. In such a case, there are nontrivial decisions made in every future period, unlike in the model with date-0 trade.

We will now, in turn, consider a series of example economies and, for each one, define equilibrium in a detailed way.

5.1.1 An endowment economy with date-0 trade

Let the economy have only one consumer with infinite life. There is no production, but the consumer is endowed with $\varpi_t \in \Re$ units of the single consumption good at each date t. Notice that the absence of a production technology implies that the consumer is unable to move consumption goods across time; he must consume all his endowment in each period, or dispose of any balance. An economy without a production technology is called an *exchange economy*, since the only economic activity (besides consumption) that agents can undertake is trading. Let the consumer's utility from any given consumption path $\{c_t\}_{t=0}^{\infty}$ be given by

$$\sum_{t=0}^{\infty} \beta^t \cdot u(c_t).$$

The allocation problem in this economy is trivial. But imagine that we deceived the consumer into making him believe that he could actually engage in transactions to buy and sell consumption goods. Then, since in truth there is no other agent who could act as his counterpart, *market clearing* would require that prices are such that the consumer is willing to have exactly ϖ_t at every t.

We can see that this requires a specific price for consumption goods at each different point in time. That is, the *commodities* here are consumption goods at different dates, and each commodity has its own price: $c_t \to p_t$. We can normalize so that $p_0 = 1$, so that prices will be relative to $t = 0$ consumption goods: a consumption good at t will cost p_t units of consumption goods at $t = 0$.

Given these prices, the value of the consumer's endowment is given by

$$\sum_{t=0}^{\infty} p_t \cdot \varpi_t.$$

The value of his expenditures is

$$\sum_{t=0}^{\infty} p_t \cdot c_t.$$

Then the budget constraint requires that

$$\sum_{t=0}^{\infty} p_t \cdot c_t \leq \sum_{t=0}^{\infty} p_t \cdot \varpi_t.$$

Notice that this assumes that trading in all *commodities* takes place at the same time: purchases and sales of consumption goods for every period are carried on at $t = 0$. This market structure is called an Arrow-Debreu-McKenzie, or date-0, market, as opposed to a sequential market structure, in which trading for each period's consumption good is undertaken in the corresponding period.

Therefore in this example, we have the following:
Definition: A Competitive Equilibrium is a vector of prices $(p_t)_{t=0}^{\infty}$ and a vector of quantities $(c_t^*)_{t=0}^{\infty}$ such that:

1. $(c_t^*)_{t=0}^{\infty} = \underset{(c_t)_{t=0}^{\infty}}{\arg\max} \left\{ \sum_{t=0}^{\infty} \beta^t \cdot u(c_t) \right\}$

$$s.t. \sum_{t=0}^{\infty} p_t \cdot c_t \leq \sum_{t=0}^{\infty} p_t \cdot \varpi_t$$
$$c_t \geq 0 \ \forall t$$

2. $c_t^* = \varpi_t \ \forall t$ (market clearing constraint)

Notice, as announced earlier, that in this trivial case, market clearing (condition 2) requires that the agent consumes exactly his endowment in each period, and this determines equilibrium prices.

Quantities are trivially determined here. Prices are not. To find the price sequence that supports the quantities as a competitive equilibrium, simply use the first-order conditions from the consumer's problem. These are

$$\beta^t u'(\varpi_t) = \lambda p_t$$

for the consumption good at t, where we have used the fact that equilibrium consumption c_t equals ϖ_t, and where λ denotes the multiplier for the budget constraint. The multiplier can be eliminated to solve for any relative price, such as

$$\frac{p_t}{p_{t+1}} = \frac{1}{\beta} \frac{u'(\varpi_t)}{u'(\varpi_{t+1})}.$$

This equation states that the relative price of today's consumption in terms of tomorrow's consumption—the definition of the (gross) real interest rate—has to equal the marginal rate of substitution between these two goods, which in this case is inversely proportional to the discount rate and to the ratio of period marginal utilities. This price is expressed in terms of primitives, and with it we have a complete solution for the competitive equilibrium for this economy.

5.1.2 The same endowment economy with sequential trade

Let's look at the same *exchange economy*, but with a *sequential markets* structure. We allow 1-period loans, which carry an interest rate of $\underbrace{R_t}_{\text{gross rate}} \equiv 1 + \underbrace{r_t}_{\text{net rate}}$, on a loan between periods $t-1$ and t.

Let a_t denote the *net asset position* of the agent at time t, that is, the net amount saved (lent) from last period.

Now we are allowing the agent to transfer wealth from one period to the next by lending 1-period loans to other agents. However, this is just a fiction as before, in the sense that, since there is only one agent in the economy, there can not actually be any loans outstanding (lending requires both a lender and a borrower!). Therefore the asset market will only clear if $a_t^* = 0 \ \forall t$: the planned net asset holding is zero for every period.

With the new market structure, the agent faces not a single, but a sequence of budget constraints. His budget constraint in period t is given by:

$$\underbrace{c_t + a_{t+1}}_{\text{uses of funds}} = \underbrace{a_t \cdot R_t^* + \varpi_t}_{\text{sources of funds}},$$

where R_t^* denotes the equilibrium interest rate that the agent takes as given. With this in hand, we have the following:
Definition: A Competitive Equilibrium is a set of *sequences* $\{c_t^*\}_{t=0}^{\infty}$, $\left\{a_{t+1}^*\right\}_{t=0}^{\infty}$, $\{R_t^*\}_{t=0}^{\infty}$ such that:

1. $\{c_t^*\}_{t=0}^{\infty} = \underset{\{c_t, a_{t+1}\}_{t=0}^{\infty}}{\arg\max} \left\{ \sum_{t=0}^{\infty} \beta^t \cdot u(c_t) \right\}$

$$s.t.\ c_t + a_{t+1} = a_t \cdot R_t^* + \varpi_t \ \forall t$$
$$c_t \geq 0 \ \forall t;\ a_0 = 0$$
$$\lim_{t \to \infty} a_{t+1} \cdot \left(\prod_{t=0}^{\infty} R_{t+1} \right)^{-1} = 0 \text{ (no-Ponzi-game condition)}$$

2. Feasibility constraint: $a_t^* = 0 \ \forall t$ (asset market clearing)

3. $c_t^* = \varpi_t \ \forall t$ (goods market clearing).

Notice that the third condition necessarily follows from the first and second ones, by Walras's law: if $n-1$ markets clear in each period, then the n^{th} one will clear as well.

To determine quantities is as trivial here (with the same result) as in the date-0 world. Prices, that is interest rates, are again available from the first-order condition for saving, the consumer's Euler equation, evaluated at $c_t^* = \varpi_t$:

$$u'(\varpi_t) = \beta \cdot u'(\varpi_{t+1}) \cdot R_{t+1}^*,$$

so that

$$R_{t+1}^* = \frac{1}{\beta} \cdot \frac{u'(\varpi_t)}{u'(\varpi_{t+1})}.$$

This expression coincides with the real interest rate in the date-0 economy, not surprisingly.

5.1.3 The neoclassical growth model with date-0 trade

Next we will look at an application of the definition of competitive equilibrium to the neoclassical growth model. We will first look at the definition of competitive equilibrium with a date-0 market structure, and then at the sequential markets structure.

The assumptions in our version of the neoclassical model are as follows:

1. The consumer is endowed with 1 unit of "time" each period, which he can allocate between labor and leisure.

2. The utility derived from the consumption and leisure stream $\{c_t,\, 1-n_t\}_{t=0}^{\infty}$ is given by

$$U\left(\{c_t,\, 1-n_t\}_{t=0}^{\infty}\right) = \sum_{t=0}^{\infty} \beta^t \cdot u\left(c_t\right).$$

   That is, we assume for the moment that leisure is not valued; equivalently, labor supply bears no utility cost. We also assume that $u(\cdot)$ is strictly increasing and strictly concave.

3. The consumer owns the capital, which he rents to firms in exchange for r_t consumption goods at t per unit of capital rented. Capital depreciates at rate δ each period.

4. The consumer rents his labor services at t to the firm for a unit rental (or wage) rate of w_t.

5. The production function of the consumption/investment good is $F(K, n)$; F is strictly increasing in each argument concave, and homogeneous of degree 1.

The following are the prices involved in this market structure:

- Consumption at every t: $c_t \to p_t$

  p_t: "within period" relative prices; if $p_0 = 1$, then p_t is the price of consumption goods at t relative to (in terms of) consumption goods at $t = 0$.

- Capital services at t: $K_t \to p_t \cdot r_t$

  r_t: rental rate; price of capital services at t relative to (in terms of) consumption goods at t.

- Labor: $n_t \rightarrow p_t \cdot w_t$

  w_t: wage rate; price of labor at t relative to (in terms of) consumption goods at t.

We thus formulate

<u>Definition</u>: A competitive equilibrium is a set of sequences:

$\{p_t^*\}_{t=0}^{\infty}$; $\{r_t^*\}_{t=0}^{\infty}$; $\{w_t^*\}_{t=0}^{\infty}$ (prices)

$\{c_t^*\}_{t=0}^{\infty}$; $\left\{K_{t+1}^*\right\}_{t=0}^{\infty}$; $\{n_t^*\}_{t=0}^{\infty}$ (quantities)

such that:

1. $\{c_t^*\}_{t=0}^{\infty}$, $\left\{K_{t+1}^*\right\}_{t=0}^{\infty}$; $\{n_t^*\}_{t=0}^{\infty}$ solve the consumer's problem:

$$\left\{c_t^*,\, K_{t+1}^*,\, n_t^*\right\}_{t=0}^{\infty} = \underset{\{c_t,\, K_{t+1},\, n_t\}_{t=0}^{\infty}}{\arg\max} \left\{\sum_{t=0}^{\infty} \beta^t \cdot u\left(c_t\right)\right\}$$

$$s.t. \sum_{t=0}^{\infty} p_t^* \cdot \left[c_t + K_{t+1}\right] = \sum_{t=0}^{\infty} p_t^* \cdot \left[r_t^* \cdot K_t + (1-\delta) \cdot K_t + 1 \cdot w_t^*\right]$$

$$c_t \geq 0 \ \forall t, \ \ k_0 \text{ given.}$$

   At every period t, capital is quoted in the same price as the consumption good. With respect to labor, recall that we have assumed that it has no utility cost; therefore $w_t \geq 0$ will imply that the consumer supplies all his time endowment to the labor market: $w_t \geq 0 \Rightarrow n_t^* = 1 \ \forall t$.

2. $\{K_t^*\}_{t=0}^{\infty}$, $\{n_t^*\}_{t=0}^{\infty}$ solve the firms' problem:

$$\forall t : (K_t^*, 1) = \underset{K_t,\, n_t}{\arg\max} \left\{p_t^* \cdot F\left(K_t, n_t\right) - p_t^* \cdot r_t^* \cdot K_t - p_t^* \cdot w_t^* \cdot n_t\right\}$$

   The firm's decision problem just involves a one-period choice—it is not of a dynamical nature (for example, we could imagine that firms live for just one period). All of the model's dynamics come from the consumer's capital accumulation problem.

   This condition may equivalently be expressed as follows:

   $\forall t : (r_t^*, w_t^*)$ satisfy:

$$\begin{aligned} r_t^* &= F_K\left(K_t^*, 1\right) \\ w_t^* &= F_n\left(K_t^*, 1\right) \end{aligned}$$

   Notice that this shows that if the production function $F\left(K, n\right)$ is increasing in n, then $n_t^* = 1$ follows.

3. Feasibility (market clearing):

$$c_t^* + K_{t+1}^* = F\left(K_t^*, 1\right) + (1-\delta) \cdot K_t^*$$

   This is known as the *one-sector* neoclassical growth model, since only one type of goods is produced, that can be used either for consumption in the current period or as capital in the following. There is also a vast literature on *multi-sector* neoclassical growth models, in which each type of physical good is produced with a different production technology, and capital accumulation is specific to each technology.

Let us now characterize the equilibrium. We first study the consumer's problem by deriving his intertemporal first-order conditions. Differentiating with respect to c_t, we obtain

$$c_t : \beta^t \cdot u'(c_t^*) = p_t^* \cdot \lambda^*,$$

where λ^* is the Lagrange multiplier corresponding to the budget constraint. Since the market structure that we have assumed consists of date-0 markets, there is only one budget and hence a unique multiplier. Consumption at $t+1$ obeys

$$c_{t+1} : \beta^{t+1} \cdot u'\left(c_{t+1}^*\right) = p_{t+1}^* \cdot \lambda^*.$$

Combining the two we arrive at

$$\frac{p_t^*}{p_{t+1}^*} = \frac{1}{\beta} \cdot \frac{u'(c_t^*)}{u'\left(c_{t+1}^*\right)}. \tag{1}$$

We can, as before, interpret $\frac{p_t^*}{p_{t+1}^*}$ as the real interest rate, and $\frac{1}{\beta} \cdot \frac{u'(c_t^*)}{u'\left(c_{t+1}^*\right)}$ as the marginal rate of substitution of consumption goods between t and $t+1$.

Differentiating with respect to capital, one sees that

$$K_{t+1} : \lambda^* \cdot p_t^* = \lambda^* \cdot p_{t+1}^* \cdot \left[r_{t+1}^* + (1-\delta)\right].$$

Therefore,

$$\frac{p_t^*}{p_{t+1}^*} = r_{t+1}^* + 1 - \delta.$$

Using Condition 2 (2'), we also find that

$$\frac{p_t^*}{p_{t+1}^*} = F_K\left(K_{t+1}^*, 1\right) + 1 - \delta. \tag{2}$$

The expression $F_K\left(K_{t+1}^*, 1\right) + (1-\delta)$ is the marginal return on capital: the marginal rate of technical substitution (transformation) between c_t and c_{t+1}. Combining expressions (1) and (2), we see that

$$u'(c_t^*) = \beta \cdot u'(c_{t+1}^*) \cdot \left[F_K\left(K_{t+1}^*, 1\right) + 1 - \delta\right]. \tag{EE}$$

Notice now that (EE) is nothing but the Euler Equation from the planner's problem! Therefore a competitive equilibrium allocation satisfies the optimality conditions for the centralized economy: the competitive equilibrium is optimal. You may recognize this as the First Welfare Theorem. We have assumed that there is a single consumer, so in this case Pareto-optimality just means utility maximization. In addition, as we shall see later, with the appropriate assumptions on $F(K,n)$ (namely, non-increasing returns to scale), an optimum can be supported as a competitive equilibrium—the Second Welfare Theorem.

5.1.4 The neoclassical growth model with sequential trade

The following are the prices involved in this market structure:

- Capital services at t: $K_t \to R_t$

  R_t: rental rate; price of capital services at t relative to (in terms of) consumption goods at t.

  Just for the sake of variety, we will now assume that R_t is the return on capital net of the depreciation costs. That is, with the notation used before, $R_t \equiv r_t + 1 - \delta$.

- Labor: $n_t \to w_t$

  w_t: wage rate; price of labor at t relative to (in terms of) consumption goods at t.

We have

Definition: A competitive equilibrium is a sequence $\left\{R_t^*,\, w_t^*,\, c_t^*,\, K_{t+1}^*,\, n_t^*\right\}_{t=0}^{\infty}$ such that:

1. $\left\{c_t^*,\, K_{t+1}^*,\, n_t^*\right\}_{t=0}^{\infty}$ solves the consumer's problem:

$$\left\{c_t^*,\, K_{t+1}^*,\, n_t^*\right\}_{t=0}^{\infty} = \underset{\{c_t,\, K_{t+1},\, n_t\}_{t=0}^{\infty}}{\arg\max} \left\{\sum_{t=0}^{\infty} \beta^t \cdot u(c_t)\right\}$$
$$s.t. \;\; c_t + K_{t+1} = K_t \cdot R_t^* + 1 \cdot w_t^*$$
$$k_0 \text{ given and a no-Ponzi-game condition.}$$

(Note that accumulating K_{t+1} is analogous to lending at t.)

2. $\left\{K_{t+1}^*,\, n_t^*\right\}_{t=0}^{\infty}$ solves the firms' problem:

$$\forall t : (K_t^*, 1) = \underset{K_t,\, n_t}{\arg\max} \left\{F(K_t, n_t) - R_t^* \cdot K_t + (1-\delta) \cdot K_t - w_t^* \cdot n_t\right\}$$

3. Market clearing (feasibility):

$$\forall t : c_t^* + K_{t+1}^* = F(K_t^*, 1) + (1-\delta) \cdot K_t^*.$$

The way that the rental rate has been presented now can be interpreted as saying that the firm manages the capital stock, funded by loans provided by the consumers. However, the capital accumulation decisions are still in the hands of the consumer (this might also be modelled in a different way, as we shall see later).

Let us solve for the equilibrium elements. As before, we build around the consumer's problem:

$$c_t : \beta^t \cdot u'(c_t^*) = \beta^t \cdot \lambda_t^*;$$

with the current market structure, the consumer faces a sequence of budget constraints, and hence a sequence of Lagrange multipliers $\{\lambda_t^*\}_{t=0}^{\infty}$. We also have

$$c_{t+1} : \beta^{t+1} \cdot u'\left(c_{t+1}^*\right) = \beta^{t+1} \cdot \lambda_{t+1}^*.$$

Then

$$\frac{\lambda_t^*}{\lambda_{t+1}^*} = \frac{u'(c_t^*)}{u'\left(c_{t+1}^*\right)}. \tag{1}$$

Differentiation with respect to capital yields

$$K_{t+1} : \beta^t \cdot \lambda_t^* = \beta^{t+1} \cdot R_{t+1}^* \cdot \lambda_{t+1}^*,$$

so that

$$\frac{\lambda_t^*}{\lambda_{t+1}^*} = \beta \cdot R_{t+1}^*. \tag{2}$$

Combining expressions (1) and (2), we obtain

$$\frac{u'(c_t^*)}{u'\left(c_{t+1}^*\right)} = \beta \cdot R_{t+1}^*. \tag{3}$$

From condition (2) of the definition of competitive equilibrium,

$$R_t^* = F_k(K_t^*, 1) + 1 - \delta. \tag{4}$$

Therefore, combining (3) and (4) we obtain:

$$u'(c_t^*) = \beta \cdot u'\left(c_{t+1}^*\right) \cdot \left[F_k(K_t^*, 1) + 1 - \delta\right]. \tag{EE}$$

This, again, is identical to the planner's Euler equation. This shows that the sequential market equilibrium is the same as the Arrow-Debreu-McKenzie date-0 equilibrium—and both are Pareto-optimal.

5.2 Recursive competitive equilibrium

Recursive competitive equilibrium uses the recursive concept of treating all maximization problems as split into decisions concerning today versus the entire future. As such, this concept thus has no room for the idea of date-0 trading: it requires sequential trading.

Instead of having sequences (or vectors), a recursive competitive equilibrium is a set of functions—quantities, utility levels, and prices, as functions of the "state": the relevant initial condition. As in dynamic programming, these functions allow us to say what will happen in the economy, and for every specific consumer, given an arbitrary choice of the initial state.

As above, we will state the definitions and the discuss their ramifications in the context of a series of examples, beginning with a treatment of the neoclassical growth model.

5.2.1 The neoclassical growth model

Let us assume again that the time endowment is equal to 1, and that leisure is not valued. Recall the central planner's problem that we analyzed before:

$$\begin{aligned} V(K) &= \max_{c,\, K' \geq 0} \{u(c) + \beta \cdot V(K')\} \\ s.t.\ c + K' &= F(K,\, 1) + (1-\delta) \cdot K \end{aligned}$$

In the decentralized recursive economy, the individual's budget constraint will no longer be expressed in terms of physical units, but in terms of sources and uses of funds at the going market prices. In the sequential formulation of the decentralized problem, these take the form of sequences of factor remunerations: $\{R_t,\, w_t\}_{t=0}^{\infty}$, with the equilibrium levels given by

$$\begin{aligned} R_t^* &= F_K(K_t^*,\, 1) + 1 - \delta \\ w_t^* &= F_n(K_t^*,\, 1)\,. \end{aligned}$$

Notice that both are a function of the (aggregate) level of capital (with aggregate labor supply normalized to 1). In dynamic programming terminology, what we have is a law of motion for factor remunerations as a function of the aggregate level of capital in the economy. If $\bar{K}$ denotes the (current) aggregate capital stock, then:

$$\begin{aligned} R &= R(\bar{K}) \\ w &= w(\bar{K}). \end{aligned}$$

Therefore, the budget constraint in the decentralized dynamic programming problem reads

$$c + K' = R(\bar{K}) \cdot K + w(\bar{K}). \tag{1}$$

The previous point implies that when making decisions, two variables are key to the agent: his own level of capital, K, and the aggregate level of capital, $\bar{K}$, which will determine his income. So the correct "syntax" for writing down the dynamic programming problem is:

$$V\left(K, \bar{K}\right) = \max_{c,\, K' \geq 0} \left\{u(c) + \beta \cdot V\left(K', \bar{K}'\right)\right\} \tag{2}$$

That is, the state variable for the consumer is $(K, \bar{K})$.

We already have the objective function that needs to be maximized and one of the restrictions, namely the budget constraint. Only $\bar{K}'$ is left to be specifieded. The economic interpretation of this is that we must determine the agent's *perceived* law of motion of aggregate capital. We assume that he will perceive this law of motion as a function of the aggregate level of capital. Furthermore, his perception will be *rational*—it will correctly correspond to the actual law of motion:

$$\bar{K}' = G(\bar{K}), \tag{3}$$

where G is a result of the economy's, that is, the representative agent's, equilibrium capital accumulation decisions.

Putting (1), (2) and (3) together, we write down the consumer's complete dynamic problem in the decentralized economy:

$$\begin{aligned} V\left(K, \bar{K}\right) &= \max_{c,\, K' \geq 0} \left\{u(c) + \beta \cdot V\left(K', \bar{K}'\right)\right\} \qquad \text{(FE)} \\ s.t.\ c + K' &= R(\bar{K}) \cdot K + w(\bar{K}) \\ \bar{K}' &= G(\bar{K}). \end{aligned}$$

(FE) is the recursive competitive equilibrium functional equation (this term includes the restrictions). The solution will yield a policy function for the individual's law of motion for capital:

$$K' = g(K, \bar{K}) = \arg\max\left\{(FE)\right\}.$$

We are already able to address our object of study:

Definition: A Recursive Competitive Equilibrium is a set of functions:

$G(\bar{K})$, $g(K, \bar{K})$ (quantities)

$V\left(K, \bar{K}\right)$ (utility level)

$R(\bar{K})$, $w(\bar{K})$ (prices)

such that:

1. $V\left(K, \bar{K}\right)$ solves (FE) and $g(K, \bar{K})$ is the associated policy function.
2. Prices are competitively determined:

$$\begin{aligned} R(\bar{K}) &= F_K(\bar{K}, 1) + 1 - \delta \\ w(\bar{K}) &= F_n(\bar{K}, 1); \end{aligned}$$

   in the recursive formulation, prices are stationary functions, rather than sequences.

3. *Consistency*:

$$G(\bar{K}) = g(\bar{K}, \bar{K})\ \forall \bar{K}$$

The third condition is the distinctive feature of the recursive formulation of competitive equilibrium. The requirement is that, whenever the individual consumer is endowed with a level of capital equal to the aggregate level (for example, only one single agent in the economy owns all the capital, or there is a measure of agents), his own individual behavior will exactly mimic the aggregate behavior. The term *consistency* points out the fact that the aggregate law of motion *perceived* by the agent must be *consistent* with the actual behavior of individuals. Consistency in the recursive framework corresponds to the idea in the sequential framework that consumers' chosen sequences of, say, capital, have to satisfy their first-order conditions given prices that are determined from firm first-order conditions *evaluated using the same sequences of capital.*

Of the three conditions defining a recursive competitive equilibrium, not one of them mentions market clearing. Will markets clear? That is, will the following equality hold?

$$\bar{c} + \bar{K}' = F(\bar{K},\, 1) + (1 - \delta) \cdot \bar{K},$$

where $\bar{c}$ denotes aggregate consumption. To answer this question, we may make use of the Euler Theorem. If the production technology exhibits constant returns to scale (that is, if the production function is homogeneous of degree 1), then that theorem delivers:

$$F(\bar{K},\, 1) + (1 - \delta) \cdot \bar{K} = R(\bar{K}) \cdot \bar{K} + w(\bar{K})$$

In economic terms, there are zero profits: the product gets exhausted in factor payment. This equation, together with the consumer's budget constraint evaluated in equilibrium ($K = \bar{K}$) implies market clearing.

Completely solving for a recursive competitive equilibrium involves more work than solving for a sequential equilibrium, since it involves solving for the functions V and g, which specify "off-equilibrium" behavior: what the agent would do if he were different from the representative agent. This calculation is important in the sense that in order to justify the equilibrium behavior we need to see that the postulated, chosen path, is not worse than any other path. $V(K, \bar{K})$ precisely allows you to evaluate the future consequences for these behavioral alternatives, thought of as one-period deviations. Implicitly this is done with the sequential approach also, although in that approach one typically simply derives the first-order (Euler) equation and imposes $K = \bar{K}$ there; knowing that the FOC is sufficient, one does not need to look explicitly at alternatives.

The known parametric cases of recursive competitive equilibria that can be solved fully include the following ones: (i) logarithmic utility (additive logarithms of consumption and leisure, if leisure is valued), Cobb-Douglas production, and 100% depreciation; (ii) isoelastic utility and linear production; and (iii) quadratic utility and linear production. It is also possible to show that, when utility is isoelastic (and no matter what form the production function takes), one obtains decision rules of the form $g(K, \bar{K}) = \lambda(\bar{K})K + \mu(\bar{K})$, where the two functions λ and μ satisfy a pair of functional equations whose solution depends on the technology and on the preference parameters. That is, the individual decision rules are linear in K, the own holdings of capital.

More in the spirit of solving for sequential equilibria, one can solve for recursive competitive equilibrium less than fully by ignoring V and g and only solve for G, using the competitive equilibrium version of the functional Euler equation. It is straightforward to show, using the envelope theorem as above in the section on dynamic programming, that this functional equation reads

$$u'\left(R(\bar{K})K + w(\bar{K}) - g(K, \bar{K})\right) = \beta u'\left(R(G(\bar{K}))g(K, \bar{K}) + w(G(\bar{K})) - g(g(K, \bar{K}), G(\bar{K}))\right)(F_1(G(K), 1) + 1 - \delta) \qquad \forall K$$

Using the Euler Theorem and consistency ($K = \bar{K}$) now lets us see that this functional equation becomes

$$u'\left(F(\bar{K}, 1) + (1-\delta)\bar{K} - G(\bar{K})\right) = \beta u'\left(F(G(\bar{K}), 1) + (1-\delta)G(\bar{K}) - G(G(\bar{K}))\right)\left(F_1(G(\bar{K}), 1) + 1 - \delta\right) \qquad \forall \bar{K},$$

which corresponds exactly to the functional Euler equation in the planning problem. We have thus shown that the recursive competitive equilibrium produces optimal behavior.

5.2.2 The endowment economy

Let the endowment process be stationary: $\varpi_t = \varpi \ \forall t$. The agent is allowed to save in the form of loans (or assets). His net asset position at the beginning of the period is given by a. Asset positions need to cancel out in the aggregate: $\overline{a} = 0$, since for every lender there must be a borrower. The definition of a recursive equilibrium is now as follows.

Definition: A Recursive Competitive Equilibrium is a set of functions

$V(a)$, $g(a)$, R

such that:

1. $V(a)$ solves the consumer's functional equation:

$$\begin{aligned} V(a) &= \max_{c \geq 0,\, a'} \{u(c) + \beta \cdot V(a')\} \\ s.t.\ c + a' &= a \cdot R + \varpi \end{aligned}$$

2. Consistency:

$$g(0) = 0$$

The consistency condition in this case takes the form of requiring the agent that has a null initial asset position to keep this null balance. Clearly, since there is a unique agent then asset market clearing requires $a = 0$. This condition will determine R as the return on assets needed to sustain this equilibrium. Notice also that R is not a function—it is a constant, since the aggregate net asset position is zero. That is, it is not correct to define a function as any $\bar{a}$ other than 0 is not feasible.

Using the functional Euler equation, which of course can be derived here as well, it is straightforward to see that R has to satisfy

$$R = \frac{1}{\beta},$$

since the u' terms cancel. This value induces agents to save zero, if they start with zero assets. Obviously, the result is the same as derived using the sequential equilibrium definition.

5.2.3 An endowment economy with 2 agents

Assume that the economy is composed by 2 agents who live forever. Agent i derives utility from a given consumption stream $\left\{c_t^i\right\}_{t=0}^{\infty}$ as per the following formula:

$$U_i\left(\left\{c_t^i\right\}_{t=0}^{\infty}\right) = \sum_{t=0}^{\infty} \beta_i^t \cdot u_i\left(c_t^i\right) \;\; i = 1, 2$$

Endowments are stationary:

$$\varpi_t^i = \varpi^i \;\forall t, \; i = 1, 2.$$

Total resource use in this economy must obey:

$$c_t^1 + c_t^2 = \varpi^1 + \varpi^2 \;\forall t.$$

Clearing of the asset market requires that:

$$\overline{a}_t \equiv a_t^1 + a_t^2 = 0 \;\forall t.$$

Notice this implies $a_t^1 = -a_t^2$; that is, at any point in time it suffices to know the asset position of one of the agents to know the asset position of the other one as well. Denote $A_1 \equiv a_1$. This is the relevant aggregate state variable in this economy (the time subscript is dropped to adjust to dynamic programming notation). Claiming that it is a state variable amounts to saying that the distribution of asset holdings will matter for prices. This claim is true except in special cases (as we shall see below), because whenever marginal propensities to save out of wealth are not the same across the two agents (either because they have different utility functions or because their common utility function makes the propensity depend on the wealth level), different prices are required to make total savings be zero, as equilibrium requires.

Finally, let q denote the current price of a one-period bond: $q = \frac{1}{R_{t,\,t+1}}$. We are now ready to state the following:

<u>Definition</u>: A Recursive Competitive Equilibrium of the 2 agent economy is a set of functions:

$g_1(a_1, A_1)\,;\; g_2(a_2, A_1)\,;\; G(A_1)$ (quantities)

$V_1(a_1, A_1)\,;\; V_2(a_2, A_1)$ (utility levels)

$q(A_1)$ (price)

such that:

1. $V_i(a_i, A_1)$ is the solution to consumer i's problem:

$$\begin{aligned} V_i(a_i, A_1) &= \max_{c^i \geq 0,\, a_i'} \left\{u_i\left(c^i\right) + \beta_i \cdot V_i\left(a_i', A_1'\right)\right\} \\ s.t.\; c^i + a_i' \cdot q(A_1) &= a_i + \varpi_i \\ A_1' &= G(A_1) \;\; \rightarrow \text{perceived law of motion for } A_1 \end{aligned}$$

The solution to this functional equation delivers the policy function $g_i(a_i, A_1)$.

2. Consistency:

$$
\begin{aligned}
G(A_1) &= g_1(A_1, A_1) \ \ \forall A_1 \\
-G(A_1) &= g_2(-A_1, A_1) \ \forall A_1.
\end{aligned}
$$

The second condition implies asset market clearing:

$$g_1(A_1, A_1) + g_2(-A_1, A_1) = G(A_1) - G(A_1) = 0.$$

Also note that q is the variable that will adjust for consistency to hold.

For this economy, it is not easy to find analytical solutions, except for special parametric assumptions. We will turn to those now. We will, in particular, consider the following question.

"AGGREGATION": Under what conditions will q be constant (that is, independent of the wealth distribution characterized by A_1)?

The answer is, one can show, that it will be constant if the following two conditions hold:

(i) $\beta_1 = \beta_2$

(ii) $u_1 = u_2$, and these utility indices belong to the Gorman class:

$$u(c) = V[f(c)]$$

where $f(\cdot)$ is affine:

$$f(c) = a + b \cdot c$$

$$V \text{ is } \begin{cases} \text{quadratic} \\ \text{exponential} \\ \text{CRRA: } \dfrac{c^{1-\sigma} - 1}{1-\sigma}. \end{cases}$$

This proposition is also valid for models with production, and can be extended to uncertainty and to the case of valued leisure. The underlying intuition is that q will be independent of the distribution of wealth if consumers' utility functions induce constant marginal propensities to save, and these propensities are identical across consumers. And it is a well-known result in microeconomics that these utility functions deliver just that.

5.2.4 Neoclassical production again, with capital accumulation by firms

Unlike in the previous examples (recall the discussion of competitive equilibrium with the sequential and recursive formulations), we will assume that firms are the ones that make capital accumulation decisions in this economy. The (single) consumer owns stock in the firm. In addition, instead of labor, we will have "land" as the second factor of production. Land will be owned by the firm.

The functions involved in this model are the dynamic programs of both the consumer and the firm:

$\bar{K}' = G(\bar{K})$	aggregate law of motion for capital
$q(\bar{K})$	current price of next period's consumption $\left(\frac{1}{\text{return on stocks}}\right)$
$V_c\left(a, \bar{K}\right)$	consumer's indirect utility as function of $\bar{K}$ and his wealth a
$g_c\left(a, \bar{K}\right)$	policy rule associated to $V_c\left(a, \bar{K}\right)$
$V_f\left(K, \bar{K}\right)$	market value (in consumption goods), of a firm with K units of initial capital, when the aggregate stock of capital is $\bar{K}$
$g_f\left(K, \bar{K}\right)$	policy rule associated to $V_f\left(a, \bar{K}\right)$.

The dynamic programs of the different agents are as follows.

1. The consumer:

$$V_c\left(a,\,\bar{K}\right) = \max_{c\geq 0,\,a'} \left\{u(c) + \beta\cdot V_c\left(a',\,\bar{K}'\right)\right\}$$

$$s.t.\ c + q\left(\bar{K}\right)\cdot a' = a$$
$$\bar{K}' = G\left(\bar{K}\right).$$

The solution to this dynamic program produces the policy rule

$$a' = g_c\left(a,\,\bar{K}\right).$$

2. The firm:

$$V_f\left(K,\,\bar{K}\right) = \max_{K'} \left\{F\left(K,\,1\right) + (1-\delta)\cdot K - K' + q\left(\bar{K}\right)\cdot V_f\left(K',\,\bar{K}'\right)\right\}$$

$$s.t.\ \bar{K}' = G\left(\bar{K}\right).$$

The solution to this dynamic program produces the policy rule

$$K' = g_f\left(K,\,\bar{K}\right).$$

We are now ready for the equilibrium definition.

Definition: A Recursive Competitive Equilibrium is a set of functions

$g_c\left(a,\,\bar{K}\right);\ g_f\left(K,\,\bar{K}\right);\ G\left(\bar{K}\right)$ (quantities)

$V_c\left(a,\,\bar{K}\right);\ V_f\left(K,\,\bar{K}\right)$ (utility level, value)

$q\left(\bar{K}\right)$ (price)

such that:

1. $V_c\left(a,\,\bar{K}\right)$ and $g_c\left(a,\,\bar{K}\right)$ are the value and policy functions, respectively, solving consumer's (FE).
2. $V_f\left(K,\,\bar{K}\right)$ and $g_f\left(K,\,\bar{K}\right)$ are the value and policy functions, respectively, solving firm's (FE).
3. Consistency 1: $g_f\left(\bar{K},\,\bar{K}\right) = G\left(\bar{K}\right)$ for all $\bar{K}$.
4. Consistency 2: $g_c\left[V_f\left(\bar{K},\,\bar{K}\right),\,\bar{K}\right] = V_f\left[G\left(\bar{K}\right),\,G\left(\bar{K}\right)\right]\ \forall\bar{K}$.

The consistency conditions can be understood as follows. The last condition requires that the consumer ends up owning 100% of the firm next period whenever he started up owning 100% of it. Notice that if the consumer starts the period owning the whole firm, then the value of a (his wealth) is equal to the market value of the firm, given by $V_f\left(\cdot\right)$. That is,

$$a = V_f\left(K,\,\bar{K}\right). \tag{1}$$

The value of the firm next period is given by

$$\text{Value of firm next period} = V_f(K',\ \bar{K}').$$

To assess this value, we need K' and $\bar{K}'$. But these come from the respective laws of motion:

$$\text{Value of firm next period} = V_f\left[g_f(K,\ \bar{K}),\,G(\bar{K})\right].$$

Now, requiring that the consumer owns 100% of the firm in the next period amounts to requiring that his desired asset accumulation, a', coincide with the value of the firm next period:

$$a' = V_f \left[g_f(K,\ \bar{K}),\, G(\bar{K}) \right].$$

But a' follows the policy rule $g_c(a,\ \bar{K})$. A substitution then gives

$$g_c(a, \bar{K}) = V_f \left[g_f(K,\ \bar{K}),\, G(\bar{K}) \right]. \tag{2}$$

Using (1) to replace a in (2), we obtain

$$g_c \left[V_f \left(K,\, \bar{K} \right), \bar{K} \right] = V_f \left[g_f(K,\ \bar{K}),\, G(\bar{K}) \right]. \tag{3}$$

The consistency condition is then imposed with $K = \bar{K}$ in (3) (and using the "Consistency 1" condition $g_f \left[\bar{K}, \bar{K} \right] = G \left[\bar{K} \right]$), yielding

$$g_c \left[V_f \left(\bar{K},\, \bar{K} \right), \bar{K} \right] = V_f \left[G(\bar{K}),\, G(\bar{K}) \right].$$

To show that the allocation resulting from this definition of equilibrium coincides with the allocation we have seen earlier (e.g., the planning allocation), one would derive functional Euler equations for both the consumer and the firm and simplify. We leave it as an exercise to verify that the outcome is indeed the optimal one.