Advanced Macro: Growth Theory Lecture 1

Marcel Fafchamps

Part I: What explains the rise in standards of living over time?

1 Neoclassical growth theory

Much of neoclassical growth theory is essentially Weberian ethics revisited: countries are more prosperous because they save more, work hard, and don't have too many children. In exchange, God's grace is visited upon all in the form of manna from heaven – technological change – that is shared equally among all people but enjoyed disproportionately by those who save and work hard. Sounds far out? Think again.

1.1 Growth accounting

Growth accounting is a method invented by Solow (1957) for decomposing growth into three (possibly more) components: accumulation of capital, increase in the labor force, and an unexplained (Solow) residual interpreted as technological change by Solow and many economists after him. The starting point is the constant returns to scale production function:

$$Q = A(t)F(K, L)$$

where L is labor and K is capital. Here output Q stands for value added, i.e., it is net of intermediate inputs. The coefficient $A(t)$ captures total factor productivity (TFP). Its variation over time is is interpreted by Solow as (Hicks-neutral) technological change. Current usage is a bit more careful, talking of the Solow residual and of total factor productivity but often refraining from crediting changes in $A(t)$ solely to technological change.

Straightforward derivation of the production function yields:

$$\dot{Q} = \dot{A}F + A\frac{\partial F}{\partial K}\dot{K} + A\frac{\partial F}{\partial L}\dot{L} \text{ or, in growth terms}$$

$$\frac{\dot{Q}}{Q} = \frac{\dot{A}}{A} + A\frac{\partial F}{\partial K}\frac{\dot{K}}{Q} + A\frac{\partial F}{\partial L}\frac{\dot{L}}{Q} \tag{1}$$

Let r, w and p be the prices of capital, labor, and the final good, respectively. Define the shares of capital and labor in value added as:

$$\begin{aligned}\omega_k &= \frac{rK}{pQ} \\ \omega_l &= \frac{wL}{pQ}\end{aligned}$$

Since production is assumed to be constant returns to scale, payments to factors of production exhaust value added. We thus have $\omega_k + \omega_l = 1$. Profit maximization further implies that

$$\begin{aligned}p\frac{\partial Q}{\partial L} &= w \\ p\frac{\partial Q}{\partial K} &= r\end{aligned}$$

Plugging into the share definitions, we get

$$\begin{aligned}\omega_k &= \frac{\partial Q}{\partial K}\frac{K}{Q} \\ \omega_l &= \frac{\partial Q}{\partial L}\frac{L}{Q}\end{aligned}$$

Replacing in (1), we get:

$$\frac{\dot{Q}}{Q} = \frac{\dot{A}}{A} + \omega_k\frac{\dot{K}}{K} + \omega_l\frac{\dot{L}}{L} \tag{2}$$

Growth $\frac{\dot{Q}}{Q}$ has been decomposed into three components: variation in the coefficient A, which is interpreted as technological change; accumulation of capital; and growth of the labor force. This formula can be further manipulated to yield a formula in terms of output per capita.

In practice, equation (2) is used to obtain an expression for the rise in TFP – or Solow residual – $\frac{\dot{A}}{A}$:

$$\frac{\dot{A}}{A} = \frac{\dot{Q}}{Q} - \omega_k\frac{\dot{K}}{K} - \omega_l\frac{\dot{L}}{L}$$

where $\frac{\dot{Q}}{Q}$ is obtained from output data, $\frac{\dot{K}}{K}$ from capital stock data, $\frac{\dot{L}}{L}$ from employment data, and share parameters ω_k and ω_l are calculated as the share of capital and labor in value added. Equation (2) thus is NOT a regression but a formula to calculate, for a country, sector or firm, TFP growth from one year to the next.

When Solow applied this formula to American growth from 1909 until 1949, he found that only 8-12.5% in growth of output per worker was attributable to capital accumulation. The 'rest', Solow argued, is due to technological change (although Solow recognized that technological change may be embodied in capital). Given this finding, you would have expected economists to focus on the sources of technological change. Instead, they spend the next thirty years refining the capital accumulation part. This is now known as neoclassical growth theory.

Before we examine this body of work in detail, it is worth mentioning that growth accounting as a methodology took a life of its own. It has been used extensively by economic historians, such as Maddison, to study the evolution of TFP over time. Macro-economists interested in business cycles have also used it to study short-term variations in productivity. It is well known, for instance, that labor productivity and total factor productivity are procyclical. Since trade cycles are not my forte, it is probably best if you learn about this from somebody else.

Growth accounting has also been used by growth economists as a descriptive tool to examine TFP differentials between countries, sectors, and over time. The TFP differentials between Europe and the US, for instance, have received much attention. Economists have also noted the fall in the rate of TFP growth after 1973-5. Obviously, the quality of inference one can draw from growth accounting exercises depends on the quality of the underlying data. Perhaps the best cautionary tale in this respect is a paper by Alwyn Young entitled 'Tale of Two Cities' and comparing Hong Kong with Singapore. The paper claimed that while the two countries had experience a very similar growth pattern, Singapore had experienced much less TFP growth than Hong Kong. This sparkled a controversy about the growth implications of this finding. It was subsequently discovered that the finding was an artifact of the way the Singapore government has over-reported the growth of its capital stock – possibly to impress investors. Once the data were corrected, the two countries were shown to have both experienced healthy TFP growth.

It is useful to note that there exists a dual equivalent to equation (2). The cost function associated with our constant returns to scale production function can be written $\frac{c(r,w)Q}{A}$ where $c(r,w)$ denotes the size-invariant unit cost function. The price=marginal cost profit maximization condition can be written:

$$p = \frac{c(r,w)}{A}$$

Taking the log derivative yields:

$$d\log p = \frac{\partial \log c(r,w)}{\partial r} d\log r + \frac{\partial \log c(r,w)}{\partial w} d\log w - d\log A$$

Since the derivative of the cost function yields the input demand, the log derivative yields the cost (or value added) share. It follows that:

$$\begin{aligned} d\log p &= \omega_k d\log r + \omega_l d\log w - d\log A \text{ which can be rewritten} \\ \frac{\dot{A}}{A} &= \omega_k \frac{\dot{r}}{r} + \omega_l \frac{\dot{w}}{w} - \frac{\dot{p}}{p} \end{aligned}$$

The above is the dual analogue of (2). If the assumptions of CRS and perfect competition are satisfied, the two formulas should give the same Solow residual. There is a literature showing, among other things, that if firms have market power the two formulas differ. This can then be construed as a test of market power. If anyone is interested, John Muellbauer has an excellent set of notes on these issues and their use in drawing inference about procyclical TFP changes over the trade cycle.

1.2 The neoclassical growth model with constant saving rate

This is the model initially pioneered by Solow in his other seminal paper of 1956. Mankiw Romer and Weil 1992 present a nicely condensed version of this growth model with constant savings rate. The basic equation is again a production function, combined with growth equations for labor and technology (exogenously given):

$$Y(t) = K(t)^{\alpha}(A(t)L(t))^{1-\alpha}$$

$$\begin{aligned} L(t) &= L(0)e^{nt} \\ A(t) &= A(0)e^{gt} \end{aligned}$$

where n is the growth rate in population/labor and g is the growth rate in technological change. Labor and technology are thus assumed to grow at exponential rates n and g. (Putting $A(t)$ in the labor aggregate is unessential since, in a Cobb-Douglas formulation, it can be factored out to look just like in the Solow model.)

The main behavioral assumption of this model is that saving is a constant proportion s of income. Solow's justification was that, in the US data he had, the savings rate was indeed constant over time.[1] Depreciation is written δ. Define $k = \frac{K}{AL}$ and $y = \frac{Y}{AL}$. We get:

$$\dot{k} = sy - (n + g + \delta)k = sk^{\alpha} - (n + g + \delta)k$$

This is the law of motion of capital. The steady state level of capital is immediately obtained by setting $\dot{k}$ to 0:

$$k^* = \left(\frac{s}{n + g + \delta}\right)^{\frac{1}{1-\alpha}}$$

From this equation, we can also define the steady state level of income per unit of effective labor $y = \frac{Y}{AL}$:

$$y^* = k^{\alpha} = \left(\frac{s}{n + g + \delta}\right)^{\frac{\alpha}{1-\alpha}}$$

---

[1] It is interesting to point out that in the Ramsey model (see below), the savings rate is constant along the balanced growth path.

It can also be shown that capital converges towards its steady state. (This should not be taken as granted: there are many dynamic models with a steady state but no path leading to it, or only certain paths leading to the steady state. We will see some later in the course.) Note: the steady state capital stock depends on the level of savings.

1.3 The neoclassical growth model with endogenous saving

Many economists felt unhappy about the assumption of a constant savings rate. Presumably, if returns to saving were higher, people ought to save more. This led to the development of the neoclassical model with endogenous saving. This model goes by many names, i.e., as the Cass-Koopmans-Ramsey-Denison model. The model eventually became the workhorse of present day macroeconomics.

I present here the simplest possible version of this model, assuming away technological change and assuming that there is a single asset k. Borrowing the notation from Intriligator, the control problem is written:

$$\max_{\{c(t)\}} W = \int_{t_0}^{\infty} e^{-\delta(t-t_0)} U(c(t))dt \text{ subject to}$$

$$\begin{aligned}
\dot{k} &= f(k) - \lambda k - c \\
k(t_0) &= k_0 \\
0 &\leq c \leq f(k)
\end{aligned}$$

where c denotes consumption (per person), k is capital (per person), δ is the rate of intertemporal time preference, λ is the depreciation rate of capital, $f(k)$ is the production function, and $U(c)$ is the utility function. Consumption is the control variable and capital is the state variable. The condition $0 \leq c \leq f(k)$ corresponds to $\{u(t)\} \in U$.

The typical treatment of the above control problem makes a whole series of assumptions to ensure that the model is well behaved. One set of assumptions is about function $f(k)$ which is called the Inada conditions:

$$\begin{aligned}
\lim_{k \to 0} f'(k) &= \infty \text{ and } f(0) > 0 \\
\lim_{k \to \infty} f'(k) &= 0
\end{aligned}$$

The first condition ensures that accumulation gets started even if $k_0 = 0$; the second condition ensures that further accumulation eventually becomes suboptimal; growth is finite. These conditions are sufficient but not necessary; weaker conditions are usually sufficient. When any of these conditions is violated, the dynamic properties of the optimal path can be very different, e.g., no accumulation at all, or eternal growth (Jones and Manuelli).

It is also customary to make a series of assumptions regarding utility, such as $U'(c) > 0$, $U''(c) < 0$, and:

$$\begin{aligned}\lim_{c\to 0} U'(c) &= \infty \\ \lim_{c\to\infty} U'(c) &= 0\end{aligned}$$

The first assumption ensures that zero consumption is never part of the optimal path. This facilitates handling the model since we do not have to worry about corner solutions and Kuhn-Tucker conditions. The second condition is not really necessary when capital depreciates. But if depreciation is 0, then combined with the second Inada condition it ensures that growth is finite – infinite accumulation is not optimal.

We are now ready to apply the maximum principle. Let $qe^{-\delta(t-t_0)}$ be the co-state variable associated with the constraint $\dot{k} = f(k) - \lambda k - c$. The Hamiltonian is:

$$H = e^{-\delta(t-t_0)}\{U(c(t)) + q[f(k) - \lambda k - c]\}$$

This variant of the Hamiltonian is a little different from the version used in the general optimal control problem in the sense that co-state is defined differently. The purpose of this transformation is to facilitate the algebra later on. The first order condition for an interior optimum is:

$$\frac{\partial H}{\partial c} = 0 \text{ yields } qe^{-\delta(t-t_0)} = e^{-\delta(t-t_0)}U'(c) \text{ or } q = U'(c)$$

The canonical equations for the co-state variable is:

$$\frac{d}{dt}e^{-\delta(t-t_0)}q(t) = -\frac{\partial H}{\partial k}$$

from which we obtain:

$$\begin{aligned}-\delta e^{-\delta(t-t_0)}q + e^{-\delta(t-t_0)}\dot{q} &= -e^{-\delta(t-t_0)}(f'(k) - \lambda)q \\ \dot{q} &= -(f'(k) - (\lambda + \delta))q\end{aligned}$$

We thus have a system of three equations with three unknowns paths, $c(t)$, $q(t)$, and $k(t)$:

$$\begin{aligned}\dot{q} &= -(f'(k) - (\lambda + \delta))q \\ \dot{k} &= f(k) - \lambda k - c \\ q &= U'(c)\end{aligned}$$

In general, this system is not easy to work with and this is about as far as we can go. Intriligator suggests a way of getting rid of the co-state variable. Totally differentiating the third equation, we get:

$$\frac{\dot{q}}{q} = \frac{U''(c)}{U'(c)}\dot{c} \equiv -\sigma(c)\frac{\dot{c}}{c}$$

where $\sigma(c)$ is the coefficient of relative risk aversion (which Intriligator calls the elasticity of marginal utility; it is also related to the intertemporal elasticity of substitution). If utility is of the CRRA form, $\sigma(c)$ is a constant. Armed with the above equation, the system boils down to two differential equations:

$$\begin{aligned}\dot{c} &= \frac{c}{\sigma(c)}(f'(k) - (\lambda + \delta)) \\ \dot{k} &= f(k) - \lambda k - c\end{aligned}$$

Now we have to switch to another part of mathematics, the solution to systems of differential equations. The first thing to do is to check that the system has a steady state. This is done by setting $\dot{c} = 0$ and $\dot{k} = 0$ and examining whether the two lines intersect. We have:

$$\begin{aligned}0 &= f'(k^*) - (\lambda + \delta) \\ 0 &= f(k^*) - \lambda k^* - c^*\end{aligned}$$

for c^* and k^* constant. The first equation determines the level of k^* such that the marginal return to capital per worker is equal to the depreciation rate plus the rate of time preference. Given k^*, the second equation determines the level of consumption as equal to output minus depreciation. The second Inada condition ensures that a steady state exists. (For a phase diagram, see Intriligator Figure 16.3.) In practical terms, this means that this economy has a resting point where it does not grow.

The next step is to check local stability. Approximating the system around the steady state with a first order Taylor approximation, we get:

$$\begin{aligned}\dot{c} &\simeq \frac{c^* f''(k^*)}{\sigma(c^*)}(k - k^*) \\ \dot{k} &\simeq \delta(k - k^*) - (c - c^*)\end{aligned}$$

(where we have made use of the fact that $f'(k^*) - \lambda = \delta$). The relevant characteristic roots are thus those of the matrix:

$$\begin{bmatrix} 0 & \frac{c^* f''(k^*)}{\sigma(c^*)} \\ -1 & \delta \end{bmatrix}$$

The characteristic polynomial has roots:

$$\frac{1}{2}\left[\delta \pm \sqrt{\delta^2 - \frac{4c^* f''(k^*)}{\sigma(c^*)}}\right]$$

Now is the part about manifolds. As expected, the roots are of opposite sign. The equilibrium point/steady state is what is called a saddle point. It has two 'manifolds' or branches, one that leads to (c^*, k^*) and one that leads away from it. The optimal path is the stable manifold because all paths that are not on the stable manifold eventually violate the transversality condition (i.e., diverge to infinity). For any given start-up value of k_0, the decision maker chooses the value of c_0 that puts the economy on the stable manifold.

We have shown that the standard neoclassical model without technological change converges to a stable steady state, which is a level of capital and income. Once this level is achieved, the economy stops growing. Growth takes place only as the economy converges to its steady state from below. This was to be expected: since by the Inada condition marginal returns to capital fall all the way to 0, it eventually becomes inefficient for an impatient representative consumer to accumulate more capital. This feature immediately makes this model an extremely unlikely candidate for explaining the growth experience of the world over the last two centuries: if anything, growth appears to have accelerated in the 20th century relative to the 19th. Growth also appears fairly constant over relatively long periods of time, while this model predict that growth would slow down over time. Finally, if countries have the same steady state, then rich countries should grow slower since they closer to their steady state. None of these predictions fit the facts.

1.4 The neoclassical model with technological change

In order to explain the stylized facts listed above, technological change is added to the neoclassical model as an exogenous process. Here I use the summary of the neoclassical model with technological change presented in Lucas (1988). The comparison between the two models is important because it brings out the concept of balanced growth path. The Lucas model also includes exogenous population growth.

The basic model starts from the familiar intertemporal optimization model for a representative agent:

$$\max \int_0^\infty e^{-\rho t} \frac{c(t)^{1-\sigma} - 1}{1-\sigma} N(t) dt \tag{3}$$

where ρ is the discount rate and $N(t)$ is (exogenously determined) population. The budget constraint (which also contains the production function) is:

$$N(t)c(t) + \dot{K}(t) = A(t)K(t)^\beta N(t)^{1-\beta} \tag{4}$$

where depreciation has been ignored for simplicity of presentation. Note that this is a model of a closed economy: local saving is invested locally. By the second welfare theory, the solution to this optimization problem (social planner problem with a single representative consumer) yields the perfect competitive equilibrium.

Given that the model is formulated in continuous time, the first step towards deriving the solution is to construct the (current-value) Hamiltonian:

$$H(K, \theta, c, t) = \frac{N}{1-\sigma}[c^{1-\sigma} - 1] + \theta[AK^\beta N^{1-\beta} - Nc]$$

where θ is a co-state variable that essentially represents the shadow price of capital. By Pontryagin's principle, the solution to the above optimization problem is to maximize the Hamiltonian at each time period. The first order conditions yield:

$$c^{-\sigma} = \theta \quad \text{(marginal utility = opportunity cost of capital)} \tag{5}$$

$$\dot{\theta} = [\rho - \beta A(t) N(t)^{1-\beta} K(t)^{\beta-1}]\theta(t) \tag{6}$$

Some algebra then follows. Lucas plugs 5 into equations 4 and 6. He gets two equations in K and θ. These two equations define a family of paths $\{K(t), \theta(t)\}$ that satisfy the $K(0)$ initial condition. The solution to the problem, or optimal path, is the unique member of the family that also satisfies the transversality condition

$$\lim_{t \to \infty} e^{-\rho t} \theta(t) K(t) = 0$$

This part is usually analyzed using a phase diagram (e.g., Aghion and Howitt page 20).

Because the above problem is hard, growth economists usually focus on a much simpler path, dubbed the balanced growth path. The balanced growth path is one possible equilibrium path for which the rates of growth of key variables are constant. Of course, being on this particular growth path requires very specific initial conditions. Why is it interesting then? Because of so-called turnpike theorems (e.g., Cass, Koopmans) that demonstrate that any economy gets closer and closer to the balanced growth path over time. It is as if the balanced growth path is a 'turnpike' (motorway): if you want to travel far, you should use the motorway. Proving turnpike theorems is hard, though. Many growth economists nevertheless use balanced growth as a benchmark in their work, even in models for which turnpike theorems do not exist... But this is not an issue in neoclassical growth theory. By the way, to get balanced growth you pretty much need to assume Cobb-Douglas production.

In the case of the above model, the growth rate that can be set to a constant is the growth rate in consumption $\frac{\dot{c}}{c} = \kappa$. As it turns out, when consumption grows as a constant rate, K and θ also do. Indeed, we immediately get $\frac{\dot{\theta}}{\theta} = -\sigma\kappa$ from equation 5. From 6 we get:

$$\beta A(t) N(t)^{1-\beta} K(t)^{\beta-1} = \rho + \sigma\kappa$$

After further algebraic manipulations (see Lucas), we get:

$$\kappa = \frac{\mu}{1-\beta}$$

$$s \equiv \frac{\dot{K}}{Nc + \dot{K}} = \frac{\beta(\kappa + \lambda)}{\rho + \sigma\kappa}$$

where $\mu \equiv \frac{\dot{A}}{A}$, λ is the (exogenous) growth rate of population, and s is the savings rate. From the above we see that the economy grows at a rate that is equal to the rate of technological progress divided by the share of labor in production.

Growth along the balanced growth path does not depend on savings, population growth, etc. It only depends on technological progress, which is assumed exogenous. This is why this body of theory is referred to as 'exogenous growth theory'.

It is interesting to compare the balanced growth path of this model to the concept of steady state of the previous sub-section. Note that if technological progress is 0, balanced growth is zero. In other words, the balanced growth path is like the steady state level of capital in a pure capital accumulation model. Balanced growth is when all transition dynamics have played out and capital accumulation rate is no longer an issue. Balanced growth is a model of long term growth for economies that have converged to their steady state capital labor ratio.

The literature then goes on to discuss the difference between level (i.e., steady state level of capital labor ratio) and growth effects (i.e., balanced growth rate). This seems to be an issue of particular interest in public finance (e.g., does taxing capital reduce growth rate or only the level of GDP without slowing growth). You may want to read Aghion and Howitt for an introduction to this literature. I do not find this literature very interesting because I do not find this model a particularly convincing model of prosperity levels across countries.

An examination of transition dynamics is quite interesting because it brings out some of the glaring inconsistencies of the neoclassical model. King and Rebelo (1993) present a series of computer simulations based on the Solow (constant savings rate) and Ramsey (endogenous savings rate) models. The results they get using standard values for labor shares, etc, indicate very slow convergence in the Solow model: 30 years for a 50% adjustment toward the steady state (balanced growth path). In contrast, they obtain very fast convergence in the Ramsey model. This is because initial rates of return on capital are extremely high (of the order of 100% and above). This induces the representative consumer to save faster at early stages of growth. Of course these results beg the question of why capital does not flow to areas/countries with such high returns to capital. The standard answer is to blame the governments of poor countries: returns to capital in poor countries are too low to attract capital or encourage local savings not because the theory is wrong but because local governments screw things up. Hence a search for the original sin: a flurry of regressions to identify the governance or institutional fault that is responsible for the failure to fit neoclassical growth theory.

1.5 Permanent growth through capital accumulation

Before we move to other models of growth, it is interesting to take a look at a particular neoclassical growth model developed by Jones and Manuelli (1990). This model is peculiar in that it manages to achieve permanent growth in a neoclassical model *without* technological change. This is of course surprising because, in the two previous sections, growth eventually stops in the absence of technological progress. The simple idea behind the paper is that, once the return on investment is equal to the depreciation rate, no growth can take place.

If, however, the return to capital forever remains above the depreciation rate, growth can go on forever.

The model of Jones and Manuelli looks very much like a Ramsey model:

$$\max \sum_{t=0}^{\infty} \beta^t u(c_t) \quad \text{subject to}$$

$$c_t + x_t = f(k_t)$$

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

where x_t stands for investment and δ is depreciation. (The authors actually have a slightly more complicated model, but it does not add to the point I wish to make.) Initial capital $k_0 \geq 0$ is given. In equilibrium,. the rental cost of capital r is such that $1 + r = 1/\beta$.

If we make the assumption that

$$\lim_{k \to \infty} f'(k) = 0$$

then there is a unique level of capital $k^* > 0$ such that $\delta k > f(k^*)$ for all $k > k^*$. Here k^* does not stand for steady state capital but for the bound on how much capital this economy can accumulate: for levels of capital above k^*, depreciation is so high that it could not be compensated even if all output was invested. In this model, capital is bounded from above; the economy cannot grow forever using only capital accumulation.

Now consider the alternative assumption:

$$f'(k) - \delta > r \ \text{ for all } k$$

This requires that the production function is of the form:

$$f(k) = ak + g(k)$$

with $a - \delta > r$ and $\lim_{k \to \infty} g'(k) = 0$. From the above it is clear that

$$\lim_{k \to \infty} f'(k) = a$$

Then we get that $\lim_{t \to \infty} c_t = \infty$: the economy grows indefinitely, and investment is forever positive.

If we consider a special case where $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$, we get the following Euler equation:

$$c_t^{-\sigma} = c_{t+1}^{-\sigma} \beta [f'(k_{t+1}) + 1 - \delta] \text{ or, after manipulation,}$$

$$\frac{c_{t+1}}{c_t} = [\beta(f'(k_{t+1}) + 1 - \delta)]^{\frac{1}{\sigma}}$$

from which we see that the growth rate declines over time but converges towards $[\beta(a + 1 - \delta)]^{\frac{1}{\sigma}}$. The limit share of capital is unity.

This model illustrates well what happens if the so-called Inada conditions are not satisfied. (The other Inada condition, $f'(0) = \infty$, guarantees that growth starts even from a zero initial level of capital.) A contrario the model shows what would happen if we could pile up more and more capital per worker without returns to capital falling to zero. Intuitively, this would require that the kind of capital that is piled up on workers varies over time, i.e., that tractors get bigger, machines get automated, etc. But this is precisely what we associate with technological progress. The model thus shows that there are alternative mathematical ways to think about the same things.

In the limit, the Jones and Manuelli model tends towards a so-called AK model, that is, a model where output can increase as a linear function of capital. In the case of the Jones and Manuelli model, AK arises because the isoquant touches one of the axes. There are other types of AK models. We will discuss AK models again when we examine human capital.

2 Appendix: The control problem

In this appendix, I introduce some of the basic methodology for the course. We begin with the general control problem, of which all growth models with a representative consumer are a special case. Understanding how the general problem is approached is more informative than limiting our attention to special cases. We begin with the continuous time version and cover the discrete time problem as a special case. There is no uncertainty. This first section borrows heavily from Intriligator.

The canonical control problem can be written:

$$\max_{\{u(t)\}} J = \int_{t_0}^{t_1} I(x,u,t)dt + F(x_1,t_1) \text{ subject to}$$

$$\dot{x} = f(x,u,t)$$

$$t_0 \text{ and } x(t_0) = x_0 \text{ given}$$

$$(x(t),t) \in T \text{ at } t = t_1$$

$$\{u(t)\} \in U$$

where $u(t)$ denotes a vector of control (i.e., decision) variables, $x(t)$ denotes a vector of state (i.e., exogenous) variables, $f(x,u,t)$ denotes the law of motion $\dot{x}$ of the state variable x (where $\dot{x} \equiv \frac{dx}{dt}$), t_0 and t_1 are initial and terminal time, $I(x,u,t)$ is called the intermediate function, $F(x_1,t_1)$ is called the final function, $f(x,u,t)$ is a vector function, $x(t_1)$ is written x_1, T is a set of values that the state variable is allowed to take at terminal time, and U is the set of values that control variables may take.

This is obviously a very general formulation, although it is possible to imagine even harder control problems, e.g., with multiple agents. At first glance it appears that the law of motion for x is restrictive as it does not allow $\dot{x}$ to depend on past values of x. This problem, however, can usually be tackled by extending the vector of state variables to allow more complicated time dependence. Infinite time problems are included in the above formulation simply by letting $t_1 = \infty$.

3 The maximum principle

One approach to solving the control problem is the maximum principle. To illustrate how it works, we follow Intriligator and begin with a slightly simplified version of the above:

$$\max_{\{u(t)\}} J = \int_{t_0}^{t_1} I(x,u,t)dt + F(x_1,t_1) \text{ subject to}$$

$$\dot{x} = f(x,u,t)$$

$$\begin{aligned} x(t_0) &= x_0 \\ x(t_1) &= x_1 \end{aligned}$$

$$\{u(t)\} \in U$$

where we also assume that functions $I(.)$, $F(.)$, and $f(.)$ are continuously differentiable functions. Setting x_1 is equivalent to setting terminal time in $x(t_1) = x_1$.

3.1 Co-state variables and Lagrangian

If we ignore boundary conditions for a moment, the above optimization problem is equivalent to maximizing a function subject to a series of constraints, which are the differential equations:

$$f(x, u, t) - \dot{x} = 0$$

Proceeding in a way analogous to the static Lagrangian problem, we create a set of new variables, called co-state variables, associated with each constraint. (Remember that $f(x, u, t)$ is a vector function, so that there may be a set of constraints at any point in time. If there is a single constraint, there is a single co-state variable.) I follow Intriligator and denote co-state variables as $y(t)$.

By analogy with the Lagrangian approach, we can form a Lagrangian expression:

$$L = \int_{t_0}^{t_1} \{I(x, u, t) + y[f(x, u, t) - \dot{x}]\}dt + F(x_1, t_1)$$

Again by analogy with the static case, a saddle point of L would yield the solution, with the difference that here the saddle point is in the space of functions.

3.2 Necessary conditions

Now consider the necessary conditions for an optimum. First consider a change in the co-state variable trajectory $y(t) + \Delta y(t)$ where $\Delta y(t)$ is a continuous function of time. This change in the co-state variable would change the value of the Lagrangian by:

$$\Delta L = \int_{t_0}^{t_1} \Delta y[f(x, u, t) - \dot{x}]dt$$

For the above to be zero for any function $\Delta y(t)$, it is necessary that:

$$f(x, u, t) - \dot{x} = 0$$

This is similar to a standard Lagrangian problem: solving for the Lagrangian multipliers yields the constraints.

To derive the other necessary conditions, it is useful to note that the term $-y\dot{x}$ can be integrated by parts to yield:

$$L = \int_{t_0}^{t_1} \{I(x,u,t) + yf(x,u,t) + \dot{y}x]\}dt + F(x_1,t_1) - [y(t_1)x(t_1) - y(t_0)x(t_0)] \quad (7)$$

For those of you who forgot where integration by parts comes from, remember that since:

$$(f(x)g(x))' = f'(x)g(x) + f(x)g'(x)$$

then:

$$f(x)g(x) = \int f'(x)g(x)dx + \int f(x)g'(x)dx$$

Rearranging and expressing the above in terms of definite integral, we obtain:

$$\int_a^b f(x)g'(x)dx = |_a^b f(x)g(x) - \int_a^b f'(x)g(x)dx$$

where is precisely what has been used above.

The first two terms under the integral in equation 7 are called the Hamiltonian function. This is a just a definition:

$$H(x,u,y,t) = I(x,u,t) + yf(x,u,t)$$

Equation 7 can thus be rewritten:

$$L = \int_{t_0}^{t_1} \{H(x,u,y,t) + \dot{y}x]\}dt + F(x_1,t_1) - [y(t_1)x(t_1) - y(t_0)x(t_0)]$$

Now consider the effect of a change in trajectory to $u(t)+\Delta u(t)$ and $x(t)+\Delta x(t)$. The change in Lagrangian becomes:

$$\Delta L = \int_{t_0}^{t_1} \{\frac{\partial H}{\partial u}\Delta u + (\frac{\partial H}{\partial x} + \dot{y})\Delta x\}dt + \left[\frac{\partial F}{\partial x_1} - y(t_1)\right]\Delta x_1$$

For a maximum, it is necessary for the change in the Lagrangian to vanish – this is how we know that we are at a saddle point. Since this must be true for any $\Delta u(t)$ and any $\Delta x(t)$, we must have:

$$\begin{aligned} \frac{\partial H}{\partial u} &= 0 \\ \frac{\partial H}{\partial x} &= -\dot{y} \\ \frac{\partial F}{\partial x_1} &= y(t_1) \end{aligned}$$

Of course, if there are no constraints, the co-state variables disappear and the necessary condition for a saddle point boils down to $\frac{\partial H}{\partial u} = 0$. Note that we have implicitly assumed that $u(t)$ is everywhere on the interior. If $u(t)$ is at a boundary, then we get Kuhn-Tucker style conditions with complementary slackness conditions. I only mention this as a clarification. Standard growth models do not hit boundaries. Another point to note is that we can write the constraint as:

$$\dot{x} = \frac{\partial H}{\partial y} (= f(x, u, t))$$

To summarize, the maximum principle can be summarized as follows:

$$u(t) = \arg\max_{u \in U} \; H(x, u, y, t) \text{ for t, } t_0 \leq t \leq t_1$$

$$\begin{aligned}
\dot{x} &= \frac{\partial H}{\partial y} \\
\dot{y} &= -\frac{\partial H}{\partial x} \\
x(t_0) &= x_0 \\
y(t_1) &= \frac{\partial F}{\partial x_1}
\end{aligned}$$

where the first equation (the maximum expression) says that u must maximize the Hamiltonian at any point in time; if u is an interior solution, then $\frac{\partial H}{\partial u} = 0$; otherwise, u is at a boundary. The four other conditions are called canonical equations. Of the four canonical equations, the first two are differential equations in state and co-state variables. The other two conditions are initial and terminal conditions, respectively. In control problems with an infinite horizon, the last set of canonical equations (terminal conditions) gets replaced by transversality conditions. The purpose of these conditions is essentially to ensure that the solution to the problem is bounded.

The usefulness of the maximum principle differs from problem to problem, a situation that is not different from the Lagrangian approach. In the overwhelming majority of cases, forming the Hamiltonian and writing down the canonical equations is totally uninformative: the problem is too complicated and the above apparatus does not help characterize the solution. This is not an anomaly: most control problems (in fact, nearly all) have no closed-form solution. Using the above to say anything about the solution is more an art than a science. Growth is no exception. In some cases, however, the solution can be characterized. To do so, one must be able to analyze systems of differential equations. To this we now turn.

4 Differential equations

Solving differential equations is not a trivial task. For the overwhelming majority of economic applications, explicitly solving the differential equations coming out of optimal control is in fact impossible because they are non-linear. Only problems that can be put in a linear or log-linear form have a closed-form solution. This is unfortunate because a closed-form solution describes the entire path of the solution, for any (reasonable) initial value. In general, we will not know the precise shape of the solution and will only be able to characterize its behavior, such as saying that it has a steady state and converges to it.

In this section, we briefly review the theory for linear differential equations, define conditions for stability, and discuss the application of these ideas to non-linear models. We borrow heavily from Luenberger, Introduction to Dynamic Systems.

4.1 Linear systems

For all intensive purposes, the only type of differential equation (or system of equations) that can easily be solved is a linear differential equation. For a single linear differential equation of the form:

$$\dot{y} = ay + b \tag{8}$$

the general solution is of the form:

$$y(t) = Ce^{at} - \frac{b}{a}$$

where C is a (yet-to-be-determined) constant that depends on the initial condition $y(0)$.

Let y^* denote the steady state of $y(t)$. This steady state – or resting state – of $y(t)$ is obtained by setting $\dot{y} = 0$ in equation 8. We get:

$$y^* = -\frac{b}{a}$$

The solution to the system shows that $y(t)$ converges to its steady state if an only if the Ce^{at} converges to 0 as $t \to \infty$, that is, iff $a < 0$. If not, the exponential term becomes arbitrarily large and the system is unbounded. What we learn from the above example is that the constant in the differential equation determines the steady state but, provided that $a \neq 0$, it plays no role in the convergence to the steady state and thus does not affect the stability of the solution. From this reasoning, we also see that C must depend on $y(0) - y^*$: if $y(0) = y^*$ then $C = 0$.

The same reasoning applies to systems of equations as well. For this reason, we consider the following system:

$$\dot{y}(t) = Ay(t) + B$$

where $y(t)$ is now a vector of variables and A and B are square matrices of constant coefficients. The steady state is:

$$y^* = -A^{-1}B$$

To facilitate analysis, we transform the above into an homogeneous system, that is, a system without constant term B. Let $z(t) \equiv y(t) - y^*$. Since y^* is a constant, $\dot{z} = \dot{y}$. We therefore obtain:

$$\begin{aligned} \dot{z}(t) &= A(z(t) + y^*) + B \\ &= Az(t) - AA^{-1}B + B \\ &= Az(t) \end{aligned}$$

If $z(t)$ converges towards 0, then $y(t)$ converges to its steady state.

Having turned the system into an homogeneous system, we note that the solution is of the form:

$$z(t) = e^{At} z(0)$$

where the expression

$$e^{At} \equiv I + At + \frac{A^2 t^2}{2!} + ... + \frac{A^k t^k}{k!} + ...$$

At first glance, the above appears extremely difficult to use even though the model is as simple as we could make it. Fortunately, things get much easier if we diagonalize the A matrix using eigen values and eigen vectors. Assume that the $n \times n$ matrix A has n distinct eigenvalues and thus n linearly independent eigenvectors. Let M be the diagonalizing matrix such that:

$$M^{-1} A M = \Lambda$$

where Λ is a diagonal matrix (i.e. a matrix with zeros off the diagonal). The Λ has the eigenvalues of A on the diagonal; the matrix M is made of the corresponding eigenvectors. Simple matrix algebra yields:

$$A^k = M \Lambda^k M^{-1}$$

The next step is to simplify expression for e^{At}. We have:

$$\begin{aligned} e^{At} &= I + At + \frac{A^2 t^2}{2!} + ... \\ &= I + M\Lambda M^{-1} t + M\Lambda^2 M^{-1} \frac{t^2}{2!} + ... \\ &= M(I + \Lambda t + \frac{\Lambda^2 t^2}{2!} + ...)M^{-1} \\ &= M e^{\Lambda t} M^{-1} \end{aligned}$$

where:

$$e^{\Lambda t} = \begin{bmatrix} e^{\lambda_1 t} & 0 & . & 0 \\ 0 & e^{\lambda_2 t} & . & . \\ . & . & . & . \\ 0 & . & . & e^{\lambda_n t} \end{bmatrix}$$

To summarize, we have shown that:

$$y(t) = (y(0) - y^*) M e^{\Lambda t} M^{-1} - A^{-1} B$$

This shows that the behavior of the entire system depends on the eigenvalues of A: if all the eigenvalues are negative, then the system converges to the steady state because all the $e^{\lambda_1 t}$ converge to 0. This illustrates that the long term behavior of the system is dominated by the largest eigen values of matrix A. The more negative the largest (closest to 0) eigenvalue is, the faster the system converges. Other eigen values determine whether the system oscillates when it approaches its resting state, or approaches it directly. (If some eigenvalues are complex, what matters is that their real part is negative.)

4.2 Non-linear systems

The above, as complicated as it is, only works for simple linear systems. What about non-linear systems? Suppose that:

$$\dot{y}(t) = G(y(t))$$

where $G(.)$ is a vector function. This problem is extremely hard and economists are like babes in the woods. Usually, economists focus on simple systems that have a steady state. Checking that a steady state exists boils down to solving a system of non-linear equations for y^*:

$$G(y^*) = 0$$

If the system only has two equations, this is equivalent to checking that the two equations intersect. If they do, there exists a steady state. Of course, it is possible for $G(y^*) = 0$ to have multiple solutions and thus for the system to have multiple steady states.

What is more difficult to show is that the system actually goes to the steady state, i.e., that the steady state is stable, either locally (things converge to the steady state if they are sufficiently close to it) or globally (things converge to the steady state irrespective of where they start). Verifying global stability turns out to be quite hard. It typically involves convexity and boundary conditions to ensure that the system remains within a given neighborhood. But much more is required to ensure that the system does not converge to a limit cycle. The theory of global stability is full of pathological cases and quite complicated. This is the world of bifurcation theory, limit cycles, chaos, and non-repeating cycles. Very few economists venture in these waters, and when they do they are seldom understood by their peers.

Verifying local stability is easier because it involves taking Taylor approximations, something economists are good at. To make the notation clearer, we index the system as follows:

$$\dot{y}_i = g_i(y^*)$$

The next step is to take a first order Taylor approximation to this system. This approximation is only valid in the immediate vicinity of the steady state, which is why it can only be used to study local stability. We get:

$$\dot{y} \simeq Jy$$

for y close to y^* and where J is the standard Jacobian matrix:

$$J \equiv \begin{bmatrix} \frac{\partial g_1(y^*)}{\partial y_1} & \frac{\partial g_1(y^*)}{\partial y_2} & . \\ \frac{\partial g_2(y^*)}{\partial y_1} & \frac{\partial g_2(y^*)}{\partial y_2} & . \\ . & . & \frac{\partial g_n(y^*)}{\partial y_n} \end{bmatrix}$$

in which each derivative is evaluated at the steady state. Stability around the steady state can then be analyzed by examining the eigen values of J. If they are all negative, the system is locally stable, that is, it converges to its steady state if it starts sufficiently close to it. In case the system has multiple steady states, the local stability of each steady state can be analyzed in the same fashion. Below we show how local stability can be checked for a growth model.

In many published articles, economists' treatment of differential equations is quite casual. Many authors simply compute the steady state and draw a phase diagram (this works only for systems of two equations, by the way). They use the equations to sign the motion of the system in each quadrant and then wave their hands. I do not like this approach because it can be very misleading. For standard growth models, stability properties are sufficiently well understood that this is not a problem in practice. But, one should be careful when constructing models that depart from the standard growth model.

4.3 Stable and unstable manifolds

The above coverage of differential equations says nothing about initial and terminal conditions. As we argued in the preceding section on optimal control, the necessary conditions for optimality include both differential equations and initial and terminal conditions. How are we supposed to combine our analysis of steady states and stability with the requirement that initial and terminal conditions be satisfied?

To clarify things, let's partition y into n_u controls u and n_x state variables x, i.e., $y = \{x, u\}$ with $n = n_u + n_x$. In additional to the differential equations, we must satisfy initial conditions on the state variables:

$$x(t_0) = x_0$$

This provides us with n_x initial conditions to tie up the differential equations. But we need an extra n_u initial conditions for the controls. Where are they to be found? By definition of optimization, the controls are free and not subject to initial conditions. Does it mean the problem is ill-defined?

No because we also have to satisfy terminal conditions. One possible case is that we must satisfy n_1 terminal conditions of the form:

$$x(t_1) = x_1$$

To satisfy these conditions, we can tinker with the initial $u(0)$ values to ensure that the n_1 targets are achieved. In general, we need as many controls as there are terminal conditions on the state variables (ignoring 'singularities' and other pathologies). In other words, we need:

$$n_u = n_1$$

The terminal conditions enable us to pick among many different paths the single path that satisfy all initial AND terminal conditions by choosing the appropriate level of initial values of u_0. The same principle applies if the terminal conditions for a finite horizon problem are replaced with transversality conditions for an infinite horizon problem.

How does the above manifest itself in the analysis of stability? In many of the growth models we will look at, we will find that the differential equations coming out of optimal control have some positive and negative roots. From what we have seen about differential equations, this seems to imply that we cannot guarantee local stability.

Normally, however, we will find that the number of negative roots is equal to the number of state variables n_x while the number of positive roots is equal to the number of terminal conditions $n_1 = n_u$. (Quite frankly I do not know what happens if $n_u \neq n_1$.) To achieve stability, initial values of the control variables have to be chosen with care so as to neutralize the positive roots. This is most easily seen if we orthogonalize/diagonalize the system $\dot{x}(t) = Ax(t)$ using:

$$x(t) \equiv Mw(t)$$

such that we obtain:

$$\begin{aligned} \dot{w}(t) &= M^{-1}\dot{x}(t) = M^{-1}AMw(t) \\ &= \Lambda w(t) \end{aligned}$$

In this case, one w variable is associated with each eigenvalue λ. 'All we have to do' to ensure that positive eigenvalues are neutralized/do not affect the stability of the system is to set to zero all initial values w_0 that correspond to positive roots. Zero w_0 remain zero forever. Since the w variables are linear transformations of the original variables of the model, this typically means selecting appropriate values of the controls u_0. In other words, if we start just right, we will still reach the steady state. For instance, if $w_0^x = \alpha x_0 - u_0 + \beta$ has a positive root, setting $u_0 = \alpha x_0 + \beta$ ensures that $w_0^x = 0$ and thus that $w^x(t) = 0$ for all t.

The set of initial values that are consistent with stability and the associate path of the system are called the stable manifold. There is also an unstable manifold that leaves the steady state and diverges indefinitely. All starting values that are not on the stable manifold diverge indefinitely (at least in linear systems). Initial values for the controls $u(0)$ must perfectly line up on the stable manifold. The reason is that values that are not on the stable manifold are not optimal because they violate the terminal conditions (or transversality conditions). For instance, suppose the two y variables are consumption and capital. Capital is the state variable x, consumption is the control variable u.

We must satisfy a single transversality condition. Given a start-up capital stock x_0, the decision maker will optimally choose consumption u_0 so as to put the economy on the stable manifold.

To summarize, the differential equations that come out of the maximum principle do not include terminal conditions. The analysis of stability based on differential equations alone therefore ignores some of the necessary conditions for optimality. These conditions must be imposed ex post by choosing initial conditions so as to eliminate paths that violate the terminal conditions.

5 References for appendix:

Intriligator, *Mathematical Optimization and Economic Theory*, Prentice Hall, 1971

Luenberger, *Introduction to Dynamic Systems: Theory, Models and Applications*, Wiley, 1979

Stokey and Lucas, *Recursive Methods in Economic Dynamics*, Harvard UP, 1989