Lecture notes by:
Matthias Kredler
Universidad Carlos III de Madrid
This version: May 2018

1 Unconstrained optimization: The basic argument used in the *Calculus of Variations*

To get acquainted with the tools used in infinite-dimensional maximization, it is useful to first have a look at a rather simple problem. Consider the profit-maximization problem of a firm that can use a continuum of labor inputs that we order on the interval $[0,1]$, given a CES production function with productivity "coefficients" $f(x)$ for each $x \in [0,1]$ and wages $w(x)$:

$$\max_{n \in C[0,1]} \overbrace{\underbrace{\left[\int_0^1 f(x)n(x)^\rho dx\right]^{1/\rho}}_{=Y(n)} - \int_0^1 n(x)w(x)dx}^{=\pi(n)},$$

where $0 < \rho < 1$, $C[0,1]$ is the space of continuous functions on $[0,1]$, and $w(x)$, $f(x)$ are given continuous functions. If this was written in sums, we would just go ahead and equalize the marginal return of a worker at position x with his wage $w(x)$. Of course, as we all know, this kind of procedure will also give us the right solution here. But let's look at the mathematically clean form of doing things.

Let's first consider what our first-order conditions for problems with finitely many arguments $x_1, \dots, x_n$ actually say: If you go into any direction in the space $\mathbb{R}^n$, the function π doesn't change its value if we consider very small changes. Since any direction in $\mathbb{R}^n$ can be represented by a linear combination of the standard basis vectors of type $(1,0,\dots,0)$, it is enough just to check the derivatives with respect to $x_1, , \dots, x_n$.

In the function space $C[0,1]$, we will now do exactly the same thing: We go into some direction (i.e. we add some element $h \in C[0,1]$ to the function n), and see what happens to the change in the profit function $\pi(n)$ if we take limits and make h very small, but still keep its basic direction. So let us consider the function

$$\tilde{\pi}(n,h,\epsilon) = \left[\int_0^1 f(x)[n(x)+\epsilon h(x)]^\rho dx\right]^{1/\rho} - \int_0^1 \big[n(x)+\epsilon h(x)\big] w(x)dx,$$

in which we have mixed in a little ϵ of the direction h. When we take the limit of very small additions and calculate the derivative $\partial\tilde{\pi}/\partial\epsilon$ and evaluate it at $\epsilon = 0$ (which we can do with standard calculus – look up the Leibniz Rule if you are not sure!) we obtain what is called the *Gateaux derivative*, which we denote it by $\delta\pi(n,h)$:[1]

$$\delta\pi(n,h) = \int_0^1 \left[f(x)\left(\frac{Y(n)}{n(x)}\right)^{1-\rho} - w(x)\right] h(x)dx = 0 \tag{1}$$

Note that this has to hold *for any* $h \in C[0,1]$ – if it didn't hold for some feasible direction h, then we should go a bit into that direction (on the increasing side, of course) and we could improve profits.

The claim is now that (1) holds for all h if and only if the term in brackets is zero for all $x \in [0,1]$, i.e.:

$$\left(\frac{Y(n)}{n(x)}\right)^{1-\rho} = w(x), \tag{2}$$

which of course is just what we guessed from the start.

The mathematical argument to prove this is the following: Suppose this was not true at some point $x' \in [0,1]$ for the optimal n^*, and say – without loss of generality – that the left-hand side (LHS) is greater than the right-hand side (RHS) in (2) at x'. Then – since n is a continuous function – we can definitely find a small interval $[x'-\varepsilon, x'+\varepsilon]$ where LHS>RHS. Now, we can construct a perturbation h that is just a little positive bump on this tiny interval. If we now look at our first-order condition (1), we see that it would be violated for this particular h. In other words, we *should actually deviate* a bit into the suggested direction h and increase profits. Hence, n^* cannot be the maximizer, a contradiction to our definition of n^*.

[1]The *Frechet derivative* is a generalization of the Gateaux derivative, but for well-behaved functions these concepts mean the same.

2 Euler equations

Now, we move on to Euler equations in continuous time. We will use the basic argument presented above, but we need one more little trick that involves integration by parts.

Consider the following problem: A planner is a given capital stock $\bar{k}_0$ at $t = 0$, and he has to hand back the economy with capital stock $\bar{k}_T$ at T to his boss. Between these two points in time, he can do whatever he wants to maximize discounted utility

$$\int_0^T e^{-\beta t} u(c_t) dt = \int_0^T e^{-\beta t} \ln(c_t) dt, \tag{3}$$

where we denote the continuous function c of time t using subscripts, i.e. c_t, to make notation nicer. The planner has to obey the following law of motion for the capital stock:

$$c_t = F(k_t) - \dot{k}_t = Ak_t^\alpha - \dot{k}_t, \tag{4}$$

where $0 < \alpha < 1$ and $\dot{k}_t$ denotes the derivative of the function k_t with respect to time (this notation is very common). This means that we're dealing with a classical production function, and there is no depreciation in the capital stock (This will make our life easier mathematically).

So the planner has to choose a function k_t on the interval $[0, T]$, where the endpoints $k_0 = \bar{k}_0$ and $k_T = \bar{k}_T$ are given. Let us use exactly the same method as before: Take some feasible deviation $h \in C[0, T]$ (h is often called a *perturbation* in the calculus of variations) from the chosen function k_t, and take the Gateaux derivative (We substitute (4) into (3) first):

$$\int_0^T e^{-\beta t} u_c(k_t, \dot{k}_t) F_k(k_t) h(t) dt = \int_0^T e^{-\beta t} u_c(k_t, \dot{k}_t) \dot{h}(t) dt \tag{5}$$

$$\int_0^T e^{-\beta t} \frac{\alpha A k_t^{\alpha-1}}{Ak_t^\alpha - \dot{k}_t} h(t) dt = \int_0^T e^{-\beta t} \frac{1}{Ak_t^\alpha - \dot{k}_t} \dot{h}(t) dt$$

Compared to before, we now have the problem that a term $\dot{h}(t)$ appears in the first-order condition. To use the same argument as before, we would like to replace this term somehow by $h(t)$.

We will use the following identity that is derived from the product rule and hence related to integration by parts:

$$f(T)g(T) - f(0)g(0) = \int_0^T \dot{f}(t)g(t) + f(t)\dot{g}(t) dt \tag{6}$$

Thus, if we restrict ourselves to perturbations h that are just tiny bumps around some interior point $t \in (0, T)$, then we have $h(0) = h(T) = 0$, and from this it follows that $\int_0^T \dot{f}(t)g(t)dt = -\int_0^T f(t)\dot{g}(t)dt$.

If we apply this insight to the first-order condition (5), we get an expression that deals with the evolution of marginal utility over time; of course, this equation is only valid for perturbations h with the property $h(0) = h(T) = 0$:

$$\int_0^T e^{-\beta t} F_k(k_t) u_c(k_t, \dot{k}_t) h_t dt = \int_0^T \frac{\partial \left[e^{-\beta t} u_c(k_t, \dot{k}_t) \right]}{\partial t} h(t) dt = 0 \tag{7}$$

On the LHS, we see the gain from increasing k_t due to higher production in t. On the right hand side, the effect is the following: Increasing k_t makes it necessary to invest *more* just before t, but *less* just after t. If marginal utility increases over time, then this leads to a loss; if marginal utility decreases, however, then we would profit from increasing k_t – it hurts us less to invest more just before t than we gain from having to invest less just after t.

Now, actually take the derivatives with respect to time to get:

$$u_c(k_t, \dot{k}_t) \left[F_k(k_t) + \beta \right] = \frac{\partial u_c(k_t, \dot{k}_t)}{\partial t} \tag{8}$$

If we use our functional assumptions on $u(\cdot)$ and $F(\cdot)$, this becomes a second-order differential equation in k. If this equation has a closed-form solution, then there is a genuine advantage of using continuous-time methods here; if not, we have solve numerically by discretizing the differential equation, and we're basically back to methods we would have used in discrete time anyway, as the shooting algorithm for example.

3 Hamilton-Jacobi-Bellman (HJB) equation: Finite horizon

Suppose we have a state variable that evolves as follows:

$$\dot{x}_t = f(x_t, h_t),$$

where x_t is the state and h_t is a control to be chosen by the planner. Also, suppose that we have the following utility function:

$$\int_0^T e^{-\beta t} u(x_t, h_t) dt$$

For now, let the utility of handing over a system with state x_T be defined by a given function $v_T(x)$.

In this case, we can do something very similar to backward induction in discrete time: Let's determine what the optimal policy is in the instant just before T having a certain state $x_{T-\varepsilon}$ in hand. Most of the time, the trade-off will be something like this: We might want to climb to a state x_T that gives us a higher final payoff $v_T(x_T)$, but this will imply some cost in $u(\cdot)$ for the time in between.

If we restrict ourselves to policies h_t that are continuous functions, then we can definitely choose ε such that h_t is basically constant on $[T-\varepsilon, T]$. For very small ε, this constant $\bar{h}$ should then approximately fulfill the following:[2]

$$v_{T-\varepsilon}(x) = \max_{\bar{h}} \left\{ \int_{T-\varepsilon}^{T} e^{-\beta t} u(x_t, \bar{h}) dt + \underbrace{\left[v_T(x) + f(x, \bar{h}) \frac{\partial v_T(x_{T-\varepsilon})}{\partial x} \varepsilon \right]}_{\text{1st-order Taylor expansion for } v_T(x_T)} \right\}$$

Now, take the value $v_T(x)$, which we can't influence by $\bar{h}$, out of the maximization and bring it to the left side. Also, notice that the integral containing u is basically taken over a constant function when ε is very small, so we take the following approximation:

$$v_{T-\varepsilon}(x) - v_T(x) = \max_{\bar{h}} \left\{ e^{-\beta T} u(x, \bar{h}) \varepsilon + f(x, \bar{h}) \frac{\partial v_T(x_{T-\varepsilon})}{\partial x} \varepsilon \right\}$$

Now we can get the optimal policy inside the brackets by balancing the pain from investing and the gain from the payoff in T. We have to set $\bar{h}$ such that

$$e^{-\beta T} \frac{\partial u(x, \bar{h})}{\partial h} + \frac{\partial f(x, \bar{h})}{\partial h} \frac{\partial v_T(x)}{\partial x} = 0 \qquad (9)$$

Since we are given $v_T(\cdot)$, x and the functional form of $u(\cdot)$ and $f(\cdot)$, this is solvable for each x.

To make headway on the evolution of the value function when we go farther to the left, we can divide by ε and take limits to get the following beautiful partial differential equation:

$$-\frac{\partial v_T(x)}{\partial t} = \max_{\bar{h}} \left\{ e^{-\beta T} u(x, \bar{h}) + f(x, \bar{h}) \frac{\partial v_T(x_{T-\varepsilon})}{\partial x} \right\} \qquad (10)$$

[2] Go over all the single statements that are implicitly used in these functions, and make sure that we are really not missing anything that is of first order!

This says that we can solve for $v_t(\cdot)$ for all $t \in [0, T]$ coming from the right: Given the derivative of v_T, for a fixed x we can get the optimal policy from the first-order condition (9) just a tiny bit left of T. To compute the new value function at this point, we can then use (10).

It is in order to make one point here about $v_t(\cdot)$: This is *not* exactly the value function that we are used to from the infinite-horizon case in the following respect: The value $v_t(\cdot)$ is *already discounted* here – to make this point clear, observe that the following would hold if the value function was indeed stationary over time, i.e. we had chosen exactly the correct stationary value function $v_T(\cdot) = v^*(\cdot)$ by chance as the ending condition:

$$\begin{aligned} v_T(x) &= e^{-\beta T} v^*(x) \\ v_t(x) &= e^{-\beta t} v^*(x) \quad \text{for all } t \in (0, T) \\ v_0(x) &= v^*(x) \end{aligned}$$

This discussion naturally leads us over to the stationary case…

4 HJB equation: Infinite horizon

The stationary value function $v^*(\cdot)$ of course has to fulfill the HJB as stated before – hence, if we just substite it into (10), equality has to hold as well:

$$\frac{\partial[e^{-\beta T} v^*(x)]}{\partial t} = \max_{\bar{h}} \left\{ e^{-\beta T} u(x, \bar{h}) + f(x, \bar{h}) \frac{\partial[e^{-\beta T} v(x)]}{\partial x} \right\}$$

Now, just taking the derivative on the left-hand side and dividing by $e^{-\beta t}$, we're already done:

$$-\beta v(x) = \max_{\bar{h}} \left\{ u(x, \bar{h}) + f(x, \bar{h}) \frac{\partial v(x)}{\partial x} \right\}$$

Note that this is only a *differential* and not a *partial differential* equation, so it is easier to study. However, we are still missing a boundary condition in order to obtain the solution. Such a boundary conditions can either come from a steady state (which we could find here by studying the Euler equation) or from explicit boundaries of the state space where the value function is given (which we don't have in this problem). Finally, we can also obtain a solution for the HJB as we do in discrete time: backward-iterate on the partial differential equation from the finite-horizon problem from the previous section until the value function converges, i.e. until the time derivative is zero for all x.