Dynamic optimization: A tool kit*

Manuel Wälti

This draft: September 2002

Contents

1 Introduction

In our microeconomics courses we have learned standard tools for maximizing an objective function (such as utility) subject to equality and inequality constraints (such as static budget constraints). Under appropriate convexity assumptions the solution to such problems usually consists of a single optimal magnitude for every choice variable, e.g. the optimal demand of a particular good, given prices and income.

In contrast, a dynamic optimization problem rises the question of what is the optimal magnitude of a choice variable in each period of time within a given time span. We could define dynamic optimization as the process of determining the paths of "control variables" and "state variables" for a dynamic system over a finite or infinite time horizon to maximize a "criterion function". There may be constraints on the final states of the system and on the 'in-flight' states and controls.

*I wish to thank Esther Brügger and Martin Wagner for helpful comments. Any errors are my own.

To illustrate this rather abstract definition let's look at an example of a dynamic optimization problem.[1] Once upon a time there was a little girl who got a cake. The girl decided to eat that cake all alone. But she was undetermined *when* she wanted to eat it. First, she thought of eating the whole cake right away. But then, nothing would be left for tomorrow and the day after tomorrow. Well, on the one hand, she thought by herself, eating cake today is better than eating it tomorrow. On the other hand, eating too much at the same time might not be the best thing to do. She imagined that the first mouthful of cake is a real treat, the second is great, the third is also nice. But the more you eat, the less you enjoy it. In the end you're almost indifferent, she thought. So, she decided to eat only a bit of the cake everyday. Then, she could eat everyday another first mouthful of cake. The girl knew that the cake would be spoiled if she kept it more than nine days. Therefore, she would eat the cake in the first ten days. Yet, how much should she eat everyday? She thought of eating everyday a piece of the same size. But if eating cake today is better than waiting for tomorrow, how can it possibly be the best to do the same today as tomorrow? If I ate just a little bit less tomorrow and a little bit more today I would be better off, she concluded. - And she would eat everyday a bit less than the previous day and the cake would last ten days long and nothing would be left in the end.

The girl's problem can be stated as follows. The girl maximizes

$$V(c_1, c_2, ..., c_T) = \sum_{t=0}^{T} \beta^t u(c_t)$$

subject to

$$k_{t+1} - k_t = -c_t$$

where

$$k_0 \text{ given}$$

$$k_{T+1} \geq 0$$

c_t is the amount of cake consumed in period t (in the context of our dynamic optimization problem, c_t represents the *control variable*). c_t yields instantaneous utility $u(c_t)$, where $u(c_t)$ satisfies the usual assumptions, i.e., $u'(\cdot) > 0$ and $u''(\cdot) < 0$. Future consumption is discounted with discount factor $0 \leq \beta \leq 1$. The present value in period 0 of the whole consumption path equals V. (The time-separable lifetime utility, V, represents the *criterion function* of our dynamic optimization problem.) T is the last day with consumption. In our story T is 9 as "today" is 0. The cake size (which represents the *state variable* in our dynamic optimization problem) is denoted by k. A first constraint on the 'in-flight' state variable requires that the cake size in period t is the previous size less the previous consumption. The original size of the cake, k_0, is given. A final constraint requires that the cake size at the terminal date must be nonnegative.

The girl's problem now is to determine the optimal path of c_t. As it stands, this problem could be solved *numerically*, e.g. with the help of the solver in Excel. However, it may also be solved *analytically*. The remainder of this handout provides you with the most relevant tools to do so, namely *optimal control* and *dynamic programming*. It is written in the style of a cookbook and

[1]The following parable is taken from Kurt Schmidheiny and Manuel Wälti, Doing economics with the computer, Session 5.

explicitly does *not* deal with the very advanced mathematics behind dynamic optimization.

This handout is mainly based on King [3], Barro and Sala-i-Martin [1], and Léonard and Van Long [2].

2 Optimal control

2.1 Discrete time

We all know the standard method of optimization with constraints, the Kuhn-Tucker Theorem.[2] To solve a dynamic optimization problem we basically apply the same theorem.

2.1.1 Finite horizon

The typical problem that we want to solve takes the following form. An agent chooses or controls a number of control variables, c_t,[3] so as to maximize an objective function subject to some constraints. These constraints are dynamic in that they describe the evolution of the state of the economy, as represented by a set of endogenous state variables which we denote by k_t. The evolution of these endogenous state variables is affected through the economic agents' choice of the control variables; they are also influenced by the variation in some exogenous state variables, x_t.

At the heart of this dynamic system are the equations describing the dynamic behavior of the states; they take the form

$$k_{t+1} - k_t = g(c_t, k_t, x_t)$$

We write these so-called *accumulation equations* as involving changes in the state variables because this eases the conversion to continuous time in our discussion below. We assume that the *initial values* of the state variables, k_0, are given. Moreover, there are *terminal conditions* on the state variables, which take the form

$$k_{T+1} \geq \bar{k}$$

The criterion (or objective) function is assumed to be a discounted sequence of flow returns, $u(c_t, k_t, x_t)$, which can represent profits, utilities, and so on. It takes the form

$$\sum_{t=0}^{T} \beta^t u(c_t, k_t, x_t)$$

This is a maximization problem with a (possibly large) number of choice variables subject to a (possibly large) number of constraints, most of which are equality constraints. Accordingly we can form the Lagrangian

$$\begin{aligned} \mathcal{L} \quad = \quad & \sum_{t=0}^{T} \beta^t u(c_t, k_t, x_t) + \sum_{t=0}^{T} \beta^t \lambda_t \left[g(c_t, k_t, x_t) + k_t - k_{t+1} \right] \\ & + \beta^{T+1} \omega_{T+1} \left[k_{T+1} - \bar{k} \right] \end{aligned}$$

[2]If not: an excellent treatment can be found in the mathematical appendix to Mas-Colell, Whinston, and Green (1995).

[3]c_t can be considered as a vector.

where λ_t denotes *current valued multipliers* (or current valued co-state variables). With the choice of *present valued multipliers*, Λ_t, the Lagrangian takes the form

$$\begin{aligned}\mathcal{L} &= \sum_{t=0}^{T} \beta^t u(c_t, k_t, x_t) + \sum_{t=0}^{T} \Lambda_t \left[g(c_t, k_t, x_t) + k_t - k_{t+1}\right] \\ &\quad + \Omega_{T+1}\left[k_{T+1} - \bar{k}\right]\end{aligned}$$

To change from one concept to the other we use the following transformation of the co-states

$$\beta^t \lambda_t = \Lambda_t$$

$$\beta^{T+1} \omega_{T+1} = \Omega_{T+1}$$

Let's look at the case of current valued multipliers and let's write $\frac{\partial u(c_t,k_t,x_t)}{\partial c_t}$ as $\frac{\partial u_t}{\partial c_t}$ etc. For the FOCs we derive the Lagrangian with respect to the controls, the states and the co-states

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 = \beta^t \frac{\partial u_t}{\partial c_t} + \beta^t \lambda_t \frac{\partial g_t}{\partial c_t} \tag{1}$$

$$\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 = -\beta^t \lambda_t + \beta^{t+1} \lambda_{t+1} \left(\frac{\partial g_{t+1}}{\partial k_{t+1}} + 1\right) + \beta^{t+1} \frac{\partial u_{t+1}}{\partial k_{t+1}} \tag{2}$$

$$\frac{\partial \mathcal{L}}{\partial k_{T+1}} = 0 = -\beta^T \lambda_T + \beta^{T+1} \omega_{T+1} \tag{3}$$

$$\frac{\partial \mathcal{L}}{\partial \left(\beta^t \lambda_t\right)} = 0 = g(c_t, k_t, x_t) + k_t - k_{t+1} \tag{4}$$

The conditions in (1) and (4) hold for $t = 0, 1, ..., T$, while those in (2) holds for $t = 0, 1, ..., T-1$. Moreover, the complementary slackness conditions have to be met:

$$\frac{\partial \mathcal{L}}{\partial \left(\beta^{T+1} \omega_{T+1}\right)} = \left[k_{T+1} - \bar{k}\right] \geq 0$$

$$\beta^{T+1} \omega_{T+1} \frac{\partial \mathcal{L}}{\partial \left(\beta^{T+1} \omega_{T+1}\right)} = \beta^{T+1} \lambda_{T+1} \left[k_{T+1} - \bar{k}\right] = 0$$

The complementary slackness conditions say that if the value of a given state variable at the terminal date is positive (i.e., $k_{T+1} \geq 0$) then its current valued shadow price must be zero. Alternatively, if its current valued shadow price at the terminal date is positive, then the agent must leave $k_{T+1} = 0$.

Application: Cake eating problem in discrete time Consider the following simplified version of the cake eating problem mentioned in Section 1. The girl chooses $c_0, c_1, ..., c_T$ that maximize

$$\textstyle\sum_{t=0}^{T} u\left(c_t\right)$$

where $u\left(c_t\right)$ satisfies $u'\left(\cdot\right) > 0$ and $u''\left(\cdot\right) < 0$, subject to

$$k_{t+1} - k_t = -c_t$$

$$k_0 \text{ given}$$

$$k_{T+1} \geq 0$$

Note that in contrast to the story above we assume here that the girl isn't impatience so that $\beta = 1$ (i.e., there is no discounting). There is one control (c_t), one endogenous state (k_t), and no exogenous state variable. Furthermore, $u(c_t, k_t, x_t) = u(c_t)$ and $g(c_t, k_t, x_t) = -c_t$. Thus, the optimality conditions are given by

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial c_t} &= 0 = \frac{\partial u_t}{\partial c_t} - \lambda_t \\
\frac{\partial \mathcal{L}}{\partial k_{t+1}} &= 0 = -\lambda_t + \lambda_{t+1} \\
\frac{\partial \mathcal{L}}{\partial k_{T+1}} &= 0 = -\lambda_T + \omega_{T+1} \\
\frac{\partial \mathcal{L}}{\partial \left(\beta^t \lambda_t\right)} &= 0 = -c_t + k_t - k_{t+1} \\
\frac{\partial \mathcal{L}}{\partial \left(\omega_{T+1}\right)} &= \left[k_{T+1} - \bar{k}\right] \geq 0 \\
\omega_{T+1} \frac{\partial \mathcal{L}}{\partial \left(\omega_{T+1}\right)} &= \lambda_{T+1} \left[k_{T+1} - \bar{k}\right] = 0
\end{aligned}$$

We end up with a set of non-linear (difference) equations. We have to solve this system choosing a value of λ_0 (or, equivalently, of c_0) so that $k_{T+1} - \bar{k} = 0$ - since $\lambda_{T+1} > 0$ due to $\frac{\partial u_{T+1}}{\partial c_{T+1}} > 0$ for any finite c_{T+1}. The most convenient way to do this is to write the above set of difference equations in the form of a nonlinear state space system (or, alternatively, as a approximated linear state space system).

2.1.2 Infinite horizon

Most models considered in economics involve economic agents with infinite planning horizons. The typical problem takes the form

$$\max_{c_t} \sum_{t=0}^{\infty} \beta^t u(c_t, k_t, x_t)$$

subject to

$$k_{t+1} - k_t = g(c_t, k_t, x_t)$$

$$k_0 \text{ given}$$

$$\lim_{t \to \infty} \lambda_t k_t = 0$$

where λ_t, as before, is the *current valued multiplier*. The terminal condition now says that k_t can be negative and grow forever in magnitude, as long as the rate of growth is *less* than λ_t; it is called the *transversality condition.*[4] Benveniste and Scheinkman (1979)[5] have shown that if

[4] An alternative form of the transversality condition, called Michel's condition, is given by $\lim_{t\to\infty} V_t = 0$.

[5] Benveniste, L.M. and J.A. Scheinkman (1979), On the Differentiability of the Value Function in Dynamic Models of Economics, Econometrica; 47(3), May, p. 727-32

- either there is discounting
- or utility is finite

then in an infinite horizon dynamic optimization problem the transversality condition is necessary.

With the choice of the current valued multipliers, the Lagrangian becomes to

$$\mathcal{L}=\sum_{t=0}^{\infty}\beta^{t}u(c_t,k_t,x_t)+\sum_{t=0}^{\infty}\beta^{t}\lambda_t\left[g(c_t,k_t,x_t)+k_t-k_{t+1}\right]$$

The optimality conditions are given by the FOCs

$$\begin{aligned}
\frac{\partial\mathcal{L}}{\partial c_t} &= 0=\beta^{t}\frac{\partial u_t}{\partial c_t}+\beta^{t}\lambda_t\frac{\partial g_t}{\partial c_t}\\
\frac{\partial\mathcal{L}}{\partial k_{t+1}} &= 0=-\beta^{t}\lambda_t+\beta^{t+1}\lambda_{t+1}\left[\frac{\partial g_{t+1}}{\partial k_{t+1}}+1\right]+\beta^{t+1}\frac{\partial u_{t+1}}{\partial k_{t+1}}\\
\frac{\partial\mathcal{L}}{\partial\left(\beta^{t}\lambda_t\right)} &= 0=g(c_t,k_t,x_t)+k_t-k_{t+1}
\end{aligned}$$

plus the initial condition and the transversality condition

$$k_0 \text{ given}$$

$$\lim_{t\to\infty}\lambda_t k_t=0$$

Application: The neoclassical growth model The benevolent social planner's problem is

$$\max\sum_{t=0}^{\infty}\beta^{t}u\left(C_t,1-N_t\right)$$

$$\text{s.t.}\quad K_{t+1}=AF\left(K_t,N_t\right)-C_t+\left(1-\delta\right)K_t$$

A solution to this problem exists because we are maximizing a continuous function over a compact set. The solution is unique since we are maximizing a concave function over a (∞-dimensional) convex set. A Lagrangian can be formed

$$L=\sum_{t=0}^{\infty}\beta^{t}u\left(C_t,1-N_t\right)+\sum_{t=0}^{\infty}\Lambda_t\left[AF\left(K_t,N_t\right)-C_t+\left(1-\delta\right)K_t-K_{t+1}\right]$$

The optimality conditions are given by the FOCs (due to the convexity of the problem these are also sufficient)

$$\begin{aligned}
\frac{\partial L}{\partial C_t} &= 0=\beta^{t}u_1\left(C_t,1-N_t\right)-\Lambda_t\\
\frac{\partial L}{\partial N_t} &= 0=\beta^{t}u_2\left(C_t,1-N_t\right)-\Lambda_t AF_2\left(K_t,N_t\right)X_t\\
\frac{\partial L}{\partial K_{t+1}} &= 0=\Lambda_t-\Lambda_{t+1}\left[AF_1\left(K_{t+1},N_{t+1}\right)+\left(1-\delta\right)\right]\\
\frac{\partial L}{\partial\Lambda_t} &= 0=AF\left(K_t,N_t\right)-C_t+\left(1-\delta\right)K_t-K_{t+1}
\end{aligned}$$

plus the boundary conditions which are given by the initial capital stock, K_0, and the transversality condition, $\lim_{t\to\infty} \Lambda_t K_{t+1} = 0$.

We end up with a set of non-linear (difference) equations. For an illustration of how such a system can be solved see e.g. Macroeconomics II, Summary 3 (Economic growth), Section 3.[6]

2.2 Continuous time

We use the Kuhn-Tucker Theorem with an infinitesimally short time period. This application is called the *Maximum Principle of Pontryagin*.

2.2.1 Finite horizon

The maximization problem now looks like

$$\max \int_0^T u\left[c\left(t\right), k\left(t\right), x\left(t\right)\right] e^{-\rho t} dt \tag{5}$$

$$\begin{aligned} \text{s.t.} \quad \frac{d}{dt} k\left(t\right) &= g\left[c\left(t\right), k\left(t\right), x\left(t\right)\right] \\ k\left(0\right) &= k_0 > 0 \text{ given} \\ k\left(T\right) &\geq \bar{k} \end{aligned}$$

As in the discrete time setting, equation (5) is called the criterion (or objective) function. The expression for $\frac{d}{dt} k\left(t\right)$ is called the accumulation (or transition) equation. Next we have the initial condition and the final constraint. For simplicity let's assume that there is just one control, one endogenous state, and one exogenous state, although a multitude of these variables could readily be included.

Digression on the discount rate in continuous time The discount factor in discrete time is β^t, where $\beta = \frac{1}{1+\rho}$. ρ denotes the rate of time preference; it expresses the impatience of an economic agent. In continuous time, β needs to be broken down to ρ since the time does not jump in units of 1 anymore. This can be done as follows

$$\beta^t = e^{t \ln \beta} = e^{(\ln \frac{1}{1+\rho})t} = e^{[\ln 1 - \ln(1+\rho)]t} = e^{[0 - \ln(1+\rho)]t} = e^{-\rho t}$$

since $\ln(1+x) \simeq x$ for small x.

The present value Hamiltonian $\mathcal{H}$ To apply the Maximum Principle we use the following cookbook procedure.[7]

1. Construct the present value Hamiltonian $\mathcal{H}$

$$\mathcal{H} = u\left[c\left(t\right), k\left(t\right), x\left(t\right)\right] e^{-\rho t} + \Lambda\left(t\right) g\left[c\left(t\right), k\left(t\right), x\left(t\right)\right]$$

where $\Lambda\left(t\right)$ again is called the present valued co-state variable or multiplier.

[6] http://www-vwi.unibe.ch/amakro/Lectures/macroii/macro.htm

[7] For a heuristic derivation see Barro and Sala-i-Martin [1].

2. Take the derivative of the Hamiltonian w.r.t. the control variable and set it to 0

$$\frac{\partial \mathcal{H}}{\partial c(t)} = 0 \tag{6}$$

3. Take the derivative of the Hamiltonian w.r.t. the state variable (the variable that appears in the differential equation above) and set it to equal the negative of the derivative of the multiplier w.r.t. time

$$\frac{\partial \mathcal{H}}{\partial k(t)} = -\frac{d}{dt}\Lambda(t) \tag{7}$$

4. Take the derivative of the Hamiltonian w.r.t. the co-state variable and set it to the derivative of the state variable w.r.t. time

$$\frac{\partial \mathcal{H}}{\partial \Lambda(t)} = \frac{d}{dt}k(t) \tag{8}$$

5. Transversality condition: Set the product of the shadow price and the state variable at the end of the planning horizon to 0[8]

$$\Lambda(T)\left[k(T) - \bar{k}\right] = 0$$

If we combine equation (6) and (7) with equation (8) (which represents nothing else than the transition equation) then we can form a *system of two differential equations* in the variables Λ and k. The final step is to find a solution to this differential equation system. For an illustrative example compare Barro and Sala-i-Martin [1], Appendix on mathematical methods, section 1.3.8.

The current value Hamiltonian $\tilde{\mathcal{H}}$

1. Construct the following current value Hamiltonian $\tilde{\mathcal{H}}$

$$\begin{aligned} \tilde{\mathcal{H}} = \mathcal{H}e^{\rho t} &= u[c(t), k(t), x(t)] + \lambda_t g[c(t), k(t), x(t)] \\ \Lambda(t) &= \lambda(t) e^{-\rho t} \end{aligned}$$

2. Take the derivative of the Hamiltonian w.r.t. the control variable and set it to 0

$$\frac{\partial \tilde{\mathcal{H}}}{\partial c(t)} = 0 \tag{9}$$

3. Take the derivative of the Hamiltonian w.r.t. the state variable (the variable that appears in the differential equation above) minus $\rho\lambda_t$ and set

[8]Generally, the transversality condition implies the following behavior for the co-state:

Case 1	$k(T) = k$	$\Rightarrow$	$\Lambda(T)$ free
Case 2	$k(T) \geq k$	$\Rightarrow$	$\Lambda(T) \geq 0$
	$k(T) > k$	$\Rightarrow$	$\Lambda(T) = 0$
Case 3	$k(T)$ free	$\Rightarrow$	$\Lambda(T) = 0$

with identical conditions for $k(0)$.

the sum of the two terms to equal the negative of the derivative of the multiplier w.r.t. time

$$\frac{\partial \tilde{\mathcal{H}}}{\partial k(t)} - \rho\lambda(t) = -\frac{d}{dt}\lambda(t) \tag{10}$$

where $\rho > 0$ (otherwise multiply it by (-1)).

4. Take the derivative of the Hamiltonian w.r.t. co-state variable and set it to the derivative of the state variable w.r.t. time

$$\frac{\partial \mathcal{H}}{\partial \lambda(t)} = \frac{d}{dt}k(t)$$

5. Transversality condition: Set the product of the shadow price and the state variable at the end of the planning horizon to 0[9]

$$\lambda(T)\left[k(T) - \bar{k}\right] = 0$$

If we combine equation (9) and (10) with the transition equation then we can form a system of two differential equations in the variables λ and k. The final step is to find a solution to this differential equation system.

Application: The cake eating problem in continuous time Consider a household who has an initial stock of cake, $k(0)$, which can be consumed over the continuous interval $0 \leq t \leq T$. The consumption of the cake generates utility $u[c(t)]$. The stock of the cake evolves through time as

$$\frac{d}{dt}k(t) = -c(t)$$

and the terminal condition is

$$k(T) \geq \bar{k}$$

where $\bar{k}$ is some positive amount of cake. Suppose that the household values cake consumption according to the utility expression

$$\int_0^T u[c(t)]\,dt$$

Note that in contrast to the story in Section 1 we assume here that the household is not impatient and, hence, $e^{-\rho t} = e^{-0t} = 1$. To solve this problem let's make use of the cookbook procedure given above. Step 1 leads to the Hamiltonian

$$\mathcal{H} = u[c(t)] - \lambda(t)c(t)$$

[9] Generally, the transversality condition implies the following behavior for the co-state:

Case 1	$k(T) = k$	$\Rightarrow$	$\Lambda(T)$ free
Case 2	$k(T) \geq k$	$\Rightarrow$	$\Lambda(T) \geq 0$
	$k(T) > k$	$\Rightarrow$	$\Lambda(T) = 0$
Case 3	$k(T)$ free	$\Rightarrow$	$\Lambda(T) = 0$

with identical conditions for $k(0)$.

Step 2 leads to the condition

$$\left(\frac{\partial \mathcal{H}}{\partial c(t)} =\right) u'[c(t)] - \lambda(t) = 0$$

Step 3 leads to the condition

$$\left(\frac{\partial \mathcal{H}}{\partial k(t)} =\right) 0 = -\dot{\lambda}(t)$$

where $\dot{\lambda}(t) = \frac{d}{dt}\lambda(t)$.

Step 4 leads to the condition

$$\left(\frac{\partial \mathcal{H}}{\partial \lambda(t)} =\right) - c(t) = \dot{k}(t)$$

Finally, the transversality condition (step 5) is given by

$$\lambda(T)\left[k(T) - \bar{k}\right] = 0$$

2.2.2 Infinite horizon with discounting

In case of an infinite horizon with discounting we can apply the same procedure as for finite horizon except that we change the transversality condition to

$$\lim_{t\to\infty} \Lambda(t)\, k(t) = 0 \quad \text{resp.} \quad \lim_{t\to\infty} \lambda(t)\, k(t) = 0$$

This means, again, that the value of the capital stock must be asymptotically 0, otherwise something valuable would be left over: If the quantity, $k(t)$, remains positive asymptotically, then the price, $\lambda(t)$, must approach 0 asymptotically. If $k(t)$ grows forever at a positive rate then the price $\lambda(t)$ must approach 0 at a faster rate so that the product, $\lambda(t)\, k(t)$, goes to 0.

Application: The neoclassical growth model with fixed labor supply Consider the following continuous time model

$$\max U = \int_0^\infty e^{-\rho t} \frac{C(t)^{1-\sigma}}{1-\sigma} dt$$

$$\begin{aligned} \text{s.t.} \quad \dot{K}(t) &= AK(t)^{1-\alpha} N^{\alpha} - C(t) - \delta K(t) \\ K(0) &= K_0 > 0 \text{ given} \end{aligned}$$

To solve this problem let's make use of our cookbook procedure. Step 1 leads to the current value Hamiltonian

$$\tilde{\mathcal{H}} = \frac{C(t)^{1-\sigma}}{1-\sigma} + \lambda(t)\left[AK(t)^{1-\alpha} N^{\alpha} - C(t) - \delta K(t)\right]$$

Step 2 leads to the condition

$$\left(\tilde{\mathcal{H}}_C =\right) C(t)^{-\sigma} - \lambda(t) = 0$$

Step 3 leads to the condition

$$\left(\tilde{\mathcal{H}}_K =\right) \lambda(t) \left[(1-\alpha) AK(t)^{-\alpha} N^{\alpha} - \delta\right] - \rho\lambda(t) = -\dot{\lambda}(t)$$

where ρ is supposed to be > 0.

Step 4 leads to the condition

$$\left(\tilde{\mathcal{H}}_\lambda =\right) AK(t)^{1-\alpha} N^{\alpha} - C(t) - \delta K(t) = \dot{K}(t)$$

Finally, the transversality condition (step 5) is given by

$$\lim_{t\to\infty} \lambda(t) k(t) = 0$$

2.3 Digression: Continuous versus discrete time

Given that the two methods are closely related, it is interesting to ask why they are both used. Continuous and discrete time differ in the following important ways:

- Phase planes: when a dynamic model is stated in continuous time one can study the qualitative dynamics (of the resulting system of differential equations) using certain graphical techniques that are not available in discrete time.
- Culture: some groups of economists learned one way and others learned another, with persisting differences.
- Solutions: whether a particular model is stated in continuous time or discrete time may lead to different solutions, i.e. there may be mathematical differences in the respective solutions.
- Closed form solutions: a closed form solution is a solution that can be arrived at by solving an equation or a set of equations, as opposed to the use of numerical methods. There are different forms of closed form solutions in discrete and continuous time. For example, in the basic growth model there is a discrete time closed form for the case with log utility and complete depreciation. In the continuous time model, there is a closed form which has other restrictions.

3 Dynamic programming

In section 2 we considered discrete optimal control theory, which is based on the familiar Kuhn-Tucker Theorem. An alternative method of solving this type of problems is the dynamic programming approach.

Recall the typical problem (notation is the same as in section 3): Find $c_0, c_1, ..., c_T$ that maximize

$$V = \sum_{t=0}^{T} \beta^t u(c_t, k_t, x_t) \tag{11}$$

subject to

$$k_{t+1} - k_t = g(c_t, k_t, x_t)$$
$$k_0 \text{ given}$$
$$k_{T+1} \geq 0$$

Dynamic programming exploits two fundamental properties of this type of problems, namely separability and additivity over time periods. More precisely,

- for any t, the functions u_t and g_t depend on t and on the state and control variables, but not on their past or future values;
- the maximand V is the sum of the net momentary utilities.

Using these two properties, Bellman (1957) enunciates an important theorem about the nature of any optimal solution of problem (11). This theorem is known as the *principle of optimality*. Roughly speaking, it says that an optimal policy has the property that at any stage t, the remaining decisions c_t^*, $c_{t+1}^*, ..., c_T^*$ must be optimal with regard to the current state k_t^*, which results from the initial state k_0 and the earlier decisions c_0^*, $c_1^*, ..., c_{t-1}^*$. This property is obviously sufficient for optimality since we require it to hold for all t: when we put $t = 1$, we have the definitions of an optimal policy. Furthermore, the property is also necessary, since any deviation from the optimal policy, even in the last period, is clearly suboptimal.

It was left to Bellman's genius to transform this rather trite, nearly tautological observation into an efficient method of solution. We now state the result formally.

3.1 Discrete time - deterministic setting

3.1.1 Finite horizon

The problem setup is given by

$$V(k_t, \delta_t, a_t) = \max_{c_t, k_{t+1}} \{u(c_t, k_t, x(\delta_t)) + \beta V(k_{t+1}, \delta_{t+1}, a_{t+1})\} \tag{12}$$
$$\text{s.t.} \quad k_{t+1} - k_t = g(c_t, k_t, x(\delta_t)) \tag{13}$$
$$a_{t+1} = a_t + 1 \tag{14}$$
$$x_t = x(\delta_t) \quad \text{with} \quad \delta_{t+1} = m(\delta_t) \tag{15}$$

(12) is the Bellman Equation, (13) is the accumulation equation, (14) is the age equation, and (15) gives the law of motion of the exogenous variable x_t as a

function of a set of exogenous state variables, δ_t, that evolve according to the - possibly nonlinear - difference equation system $\delta_{t+1} = m(\delta_t)$. For an example compare the application of stochastic dynamic programming below.

Note that we have converted the many-period optimization problem given above into a two period optimization problem, which involves trading off between the current return $u(c_t, k_t, x(\delta_t))$ and the future value $V(k_{t+1}, \delta_{t+1}, a_{t+1})$.

To solve the problem, we begin at the terminal value and proceed by backward induction. This process is frequently called value iteration, as it involves taking initial value function $V(k_{t+1}, \delta_{t+1}, a_{t+1})$, finding the optimal level of the right hand side of the Bellman equation at each k_t, δ_t and thereby constructing a new value function $V(k_t, \delta_t, a)$. We can also write now

$$\begin{aligned}
V(k_t, \delta, a) &= \max_{c,k'} \{u(c, k, x(\delta)) + \beta V(k', \delta', a')\} \\
\text{s.t.} \quad k' - k &= g(c, k, x(\delta)) \\
a' &= a + 1 \\
\delta' &= m(\delta)
\end{aligned}$$

On the Bellman Equation, we now apply the standard method:

$$\mathcal{L} = u(c, k, x(\delta)) + \beta V(k', \delta', a') + \lambda \left[g(c, k, x(\delta)) + k - k'\right]$$

where λ_t denotes the shadow price. The FOCs are

$$\begin{aligned}
\frac{\partial u(c, k, x(\delta))}{\partial c} + \lambda \frac{\partial g(c, k, x(\delta))}{\partial c} &= 0 \\
-\lambda + \beta \frac{\partial V(k', \delta', a')}{\partial k'} &= 0 \\
g(c, k, x(\delta)) + k - k' &= 0
\end{aligned}$$

The first FOC then is the derivation of $\mathcal{L}$ with respect to the control, the second with respect to the state at time $t+1$, and the third with respect to the Lagrange multiplier yielding the constraint of state accumulation. To get an expression for $\frac{\partial V(k', \delta', a')}{\partial k'}$ we need the envelope theorem (for a rigorous exposition compare the relevant literature)

$$\frac{\partial V(k, \delta, a)}{\partial k} \equiv \frac{\partial u(c, k, x(\delta)}{\partial k} + \lambda [\frac{g(c, k, x(\delta))}{\partial k} + 1]$$

Before plugging in, we change the subscripts since we need $\frac{\partial V(k,\delta)}{\partial k}$ of the subsequent period.

3.1.2 Infinite horizon

In an infinite time horizon setting, we replace the terminal condition in the typical problem above by the transversality condition

$$\lim_{t \to \infty} \lambda_t k_t = 0$$

Apart from this new condition, infinite horizon optimization is not different from the finite horizon case. Bellman's principle of optimality at once tells us why. Consider any finite horizon subproblem with the initial and terminal conditions fixed by the larger problem For the subproblem, the maximum principle conditions apply. But the initial and terminal times of the subproblem could be arbitrary, so the conditions must in fact hold for the entire range $(0, \infty)$.

Application: Non-stochastic dynamic programming We consider a non-stochastic baseline model of investment where firms face a perfectly elastic supply of capital goods and can adjust their capital stocks costlessly. Suppose that the profits of the firm can be written as

$$p(\mu_t) f(k_t) - \rho(\mu_t) i_t \equiv u(i_t, k_t, \mu_t)$$

where $p(\mu_t)$ may be interpreted as an output price or a productivity shock; $f(k_t)$ is a positive, increasing and strictly concave production function (i.e. we abstract from labor within the scope of this model); $\rho(\mu_t)$ is the investment good price; and i_t is the quantity of investment expenditure. The firm's capital accumulation will be described by

$$k_{t+1} - k_t = i_t - d \cdot k_t \equiv g(i_t, k_t)$$

The Bellman equation for the infinite horizon problem is

$$V(k_t, \mu_t) = \max_{i_t, k_{t+1}} \{u(i_t, k_t, \mu_t) + \beta V(k_{t+1}, \mu_{t+1})\}$$

where maximization takes place subject to $k_{t+1} - k_t = g(i_t, k_t)$ and the dynamic equations for the exogenous states follow $\delta_{t+1} = m(\mu_t)$.

Since this is a constrained optimization problem, we may form the Lagrangian

$$L = \{u(i_t, k_t, \mu_t) + \beta V(k_{t+1}, \mu_{t+1})\} + \lambda_t [g(i_t, k_t) - k_{t+1} + k_t]$$

The FOCs are

$$\begin{aligned} i_t &: \quad 0 = \frac{\partial u(i_t, k_t, \mu_t)}{\partial i_t} + \lambda_t \frac{\partial g(i_t, k_t)}{\partial i_t} \\ k_{t+1} &: \quad 0 = -\lambda_t + \beta \frac{\partial V(k_{t+1}, \mu_{t+1})}{\partial k_{t+1}} \\ \lambda_t &: \quad 0 = g(i_t, k_t) - k_{t+1} + k_t \end{aligned}$$

Specifically, the first FOC is $-\rho(\mu_t) + \lambda_t = 0$ or, equivalently, $\lambda_t = \rho(\mu_t)$. To get an expression for $\frac{\partial V(k_{t+1}, \mu_{t+1})}{\partial k_{t+1}}$ we need the envelope theorem

$$\frac{\partial V(k_t, \mu_t)}{\partial k_t} = \frac{\partial u(i_t, k_t, \mu_t)}{\partial k_t} + \lambda_t \left[\frac{\partial g(i_t, k_t)}{\partial k_t} + 1\right]$$

where the second term on the RHS is the derivative of the budget constraint with respect to k_t. $\frac{\partial V(k_{t+1}, \mu_{t+1})}{\partial k_{t+1}}$ is computed by updating the resulting expression to $t+1$. This yields

$$\frac{\partial V(k_{t+1}, \mu_{t+1})}{\partial k_{t+1}} = p(\mu_{t+1}) \frac{\partial f(k_{t+1})}{\partial k_{t+1}} + \lambda_{t+1}(1-d)$$

Application: Stochastic dynamic programming Dynamic programming is particularly well suited to optimization problems that combine time and uncertainty.

Consider a representative agent with preferences over consumption, c, and leisure, l, that are represented by the expected utility function,

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, l_t)$$

The momentary utility, $u(c_t, l_t)$, satisfies standard assumptions.

The representative agent faces three constraints at each date. The first says that the total uses of output for consumption and investment do not exceed total output, which is produced from capital and labor via a production function that is shifted by productivity shocks, a_t

$$c_t + i_t = a_t f(k_t, n_t)$$

The second says that labor plus leisure is equal to the time endowment (normalized to unity)

$$n_t + l_t = 1$$

The third says that the future capital stock, k_{t+1}, evolves according to the net result of investment and depreciation

$$k_{t+1} - k_t = i_t - dk_t$$

Note that the capital stock at date t, k_t, is a predetermined variable, i.e. the result of prior investment decisions.

To study the representative agent's optimality decisions, it is necessary to be explicit about how the exogenous variable, a_t, evolves through time. The general approach is to assume that a_t is a function of a variable μ_t, i.e. $a_t = a(\mu_t)$. We also assume that μ_t is a *Markov process*, so that the conditional distribution of μ_{t+1} depends only on μ_t and not on any additional past history. We call μ_t the exogenous state variable of the model. For instance, assume that $a_t = a_{t-1}^{\rho}\mu_t$, where μ_t is white noise.

We are now ready to set up the dynamic program solved by the representative agent:

$$V(k_t, \mu_t) = \max_{i_t, k_{t+1}} \{u(c_t, l_t) + \beta E_t V(k_{t+1}, \mu_{t+1} | \mu_t)\}$$

subject to

$$\begin{aligned}
\lambda_{1,t} &: \quad 0 = a_t f(k_t, n_t) - c_t - i_t \equiv g_1[c_t, i_t, n_t, k_t, a(\mu_t)] \\
\lambda_{2,t} &: \quad 0 = 1 - n_t - l_t \equiv g_2[n_t, l_t] \\
\lambda_{3,t} &: \quad k_{t+1} - k_t = i_t - dk_t \equiv g_3[i_t, k_t]
\end{aligned}$$

Since this is a constrained optimization problem, we may form the Lagrangian

$$\begin{aligned}
\mathcal{L} = \; & \{u(c_t, l_t) + \beta E_t V(k_{t+1}, \mu_{t+1})\} \\
& + \lambda_{1,t}[a_t f(k_t, n_t) - c_t - i_t] \\
& + \lambda_{2,t}[1 - n_t - l_t] \\
& + \lambda_{3,t}[i_t - dk_t - k_{t+1} + k_t]
\end{aligned}$$

The derivation of the FOCs for the two control variables c_t and l_t raises no problems

$$
\begin{aligned}
c_t &: \quad 0 = \frac{\partial u\left(c_t, l_t\right)}{\partial c_t} - \lambda_{1,t} \\
l_t &: \quad 0 = \frac{\partial u\left(c_t, l_t\right)}{\partial l_t} - \lambda_{2,t}
\end{aligned}
$$

The FOC for the control variable n_t is derived as follows. Recall that *in general* the FOC for a control variable is given by[10]

$$\frac{\partial u(c_t, k_t, x_t)}{\partial c_t} + \lambda \frac{\partial g(c_t, k_t, x_t)}{\partial c_t} = 0$$

In the case at hand, the control variable n_t does not show up in the momentary utility of the representative agent, $u\left(c_t, l_t\right)$. Moreover, n_t appears in accumulation equation $g_1\left[c_t, i_t, n_t, k_t, a\left(\mu_t\right)\right]$ and $g_2\left[n_t, l_t\right]$. It follows that the FOC is given by

$$0 + \lambda_{1,t} \frac{\partial g_1\left[c_t, i_t, n_t, k_t, a\left(\mu_t\right)\right]}{\partial n_t} + \lambda_{2,t} \frac{\partial g_2\left[n_t, l_t\right]}{\partial n_t} = 0$$

or, more specific,

$$n_t : 0 = \lambda_{1,t} \frac{a_t f\left(k_t, n_t\right)}{\partial n_t} - \lambda_{2,t}$$

A similar logic applies to the control variable i_t, which appears in accumulation equation $g_1\left[c_t, i_t, n_t, k_t, a\left(\mu_t\right)\right]$ and $g_3\left[i_t, k_t\right]$ (but not in the momentary utility function). The FOC is given by

$$i_t : 0 = -\lambda_{1,t} + \lambda_{3,t}$$

In general the FOC for a state variable is given by[11]

$$-\lambda_t + \beta E_t \left\{ \frac{\partial V\left(k_{t+1}, \mu_{t+1}\right)}{\partial k_{t+1}} \right\} = 0$$

To get an expression for $\frac{\partial V\left(k_{t+1}, \mu_{t+1}\right)}{\partial k_{t+1}}$ we need the envelope theorem:

$$\frac{\partial V(k_t, \delta_t)}{\partial k_t} \equiv \frac{\partial u\left(c_t, k_t, x_t\right)}{\partial k_t} + \lambda_t\left[\frac{g(c_t, k_t, x_t)}{\partial k_t} + 1\right]$$

In the case at hand, the state variable k_t does not show up in the momentary utility function of the representative household, $u\left(c_t, l_t\right)$. Moreover, k_t appears in accumulation equation $g_1\left[c_t, i_t, n_t, k_t, a\left(\mu_t\right)\right]$ and $g_3\left[i_t, k_t\right]$. Thus,

$$\frac{\partial V(k_t, \delta_t)}{\partial k_t} = 0 + \lambda_{1,t} \left[\frac{g_1\left[c_t, i_t, n_t, k_t, a\left(\mu_t\right)\right]}{\partial k_t} \right] + \lambda_{3,t} \left[\frac{g_3\left[i_t, k_t\right]}{\partial k_t} + 1 \right]$$

[10] Be aware of the following point: The general function $u(c_t, k_t, x_t(\delta_t))$ and the momentary utility function of the problem at hand, $u\left(c_t, l_t\right)$, use the same notation. Also, in the general problem c_t stands for control variables and k_t stands for state variables, whereas in the problem at hand c_t denotes consumption (just one of several control variables) and k_t denotes physical capital (the only endogenous state variable).

[11] Compare the previous footnote.

Why has the term +1 been skipped in the expression $\lambda_{1,t}\left[\frac{g_1[c_t,i_t,n_t,k_t,a(\mu_t)]}{\partial k_t}\right]$? Well, as you can see above, k_t does not show up on the LHS of accumulation equation 1. Changing the subscripts and substituting yields

$$k_{t+1} : 0 = -\lambda_t + \beta E_t \left\{ \lambda_{1,t+1} \left[\frac{a_{t+1} f\left(k_{t+1}, n_{t+1}\right)}{\partial k_{t+1}} \right] + \lambda_{3,t+1}\left(1 - d\right) \right\}$$

The last three FOCs are

$$\begin{aligned} \lambda_{1,t+1} &: \quad 0 = a_t f\left(k_t, n_t\right) - c_t - i_t \\ \lambda_{2,t+1} &: \quad 0 = 1 - n_t - l_t \\ \lambda_{3,t+1} &: \quad 0 = i_t - dk_t - k_{t+1} + k_t \end{aligned}$$

3.2 Continuous time

To be written.

References

[1] Barro, R.J. and X. Sala-i-Martin (1995), Economic growth, McGraw-Hill

[2] Léonard, D. and N. Van Long (1992), Optimal control theory and static optimization in economics, Cambridge University Press

[3] King, R.G. (?), Notes on dynamic optimization, handout