Optimal control and dynamic programming

Paul Schrimpf

UBC
Economics 526

November 6, 2013

1. Introduction
2. Differentiation in vector spaces
3. Optimization in vector spaces
4. Optimal control
   - Continous time optimal control
   - Application: optimal contracting with a continuum of types
   - Discrete time optimal control
5. Dynamic programming

Optimal control and dynamic programming

Paul Schrimpf

Introduction

Differentiation in vector spaces

Optimization in vector spaces

Optimal control

Continous time optimal control

Application: optimal contracting with a continuum of types

Discrete time optimal control

Dynamic programming

Section 1

Introduction

Example: consumption and savings

- Infinite horizon consumption-savings problem,

$$\max_{\{c_t\}_{t=0}^{\infty},\{s_t\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t) \text{ s.t. } s_{t+1} = (1 + r_t)(s_t - c_t),$$

- Countably infinite c_t and s_t

Example: contracting with continuum of types

- Type θ gets 0 utility from not buying the good, and $\theta\nu(q) - T$ from buying q units of the good at cost T
- $\theta \in [\theta_l, \theta_h]$, density f_θ
- Menu of contracts $(q(\theta), T(\theta))$ such that type θ will choose contract $(q(\theta), T(\theta))$

$$\max_{q(\theta),T(\theta)} \int_{\theta_l}^{\theta_h} \left[T(\theta) - cq(\theta)\right] f_\theta(\theta) d\theta$$

s.t.

$$\theta\nu\left(q(\theta)\right) - T(\theta) \geq 0 \forall\theta \tag{1}$$

$$\theta\nu\left(q(\theta)\right) - T(\theta) \geq \max_{\tilde{\theta}} \theta\nu\left(q(\tilde{\theta})\right) - T(\tilde{\theta}) \forall\theta \tag{2}$$

- These problems can be solved using same techniques as before
  - First and second order conditions
  - Involves differentiation in infinite dimensional vector spaces
  - This approach sometimes called “calculus of variations”
- Optimal control and dynamic programming are special tricks to make solving these problems easier
  - Can be derived from calculus of variations

Optimal control and dynamic programming

Paul Schrimpf

Introduction

Differentiation in vector spaces

Optimization in vector spaces

Optimal control

Continous time optimal control

Application: optimal contracting with a continuum of types

Discrete time optimal control

Dynamic programming

Section 2

Differentiation in vector spaces

Differentiation in vector spaces

$$\max_{x \in U} f(x) \text{ s.t } h(x) = c$$
$$g(x) \leq b$$

- $x \in U \subseteq V$, a Banach space
  - Countably infinite problems, usually

$$V = \ell^p = \{\{x_t\}_{t=1}^\infty : \left(\sum_{t=1}^\infty |x_t|^p\right)^{1/p} < \infty\}$$

  - Uncountably infinite problems, usually $V = \mathcal{L}^p$ or $V = C^k$
- $f : V \to \mathbb{R}$, derivative $Df_x : V \to \mathbb{R}$ is continuous and linear such that

$$\lim_{\|h\| \to 0} \frac{\|f(x+h) - f(x) - Df_x h\|}{\|h\|} = 0$$

Definition

Let V and W be Banach spaces and $f : V \to W$. The **Fréchet derivative** of f at $x \in V$ is a continuous linear transformation from V to W, denoted Df_x such that

$$\lim_{\|h\| \to 0} \frac{\|f(x+h) - f(x) - Df_x h\|}{\|h\|} = 0.$$

Definition

The **directional (Gâteaux) derivative** in direction v at x is

$$df(x;v) = \lim_{\alpha \to 0} \frac{f(x+\alpha v) - f(x)}{\alpha}.$$

where $\alpha \in \mathbb{R}$ is a scalar.

Lemma

If $f : V \to W$ is Fréchet differentiable at x, then the Gâteaux derivative, $df(x; v)$, exists for all $v \in V$, and

$$df(x; v) = Df_x v.$$

Lemma

If $f : V \to W$ has Gâteaux derivatives that are linear in v and "continuous" in x in the sense that $\forall \epsilon > 0\ \exists \delta > 0$ such that if $\|x_1 - x\| < \delta$, then

$$\sup_{v \in V} \frac{\|df(x_1; v) - df(x; v)\|}{\|v\|} < \epsilon$$

then f is Fréchet differentiable with $Df_{x_0} v = df(x; v)$.

Dual spaces

Lagrange multipliers are elements of dual spaces

Definition

Let V be a vector space. The **dual space** of V, denote V^* is the set of all (continuous) linear functionals, $v^* : V \to \mathbb{R}$.

Examples

- $(\mathbb{R}^n)^* = \mathbb{R}^n$
- $(\ell^p)^* = \ell^q$ where $\frac{1}{p} + \frac{1}{q} = 1$
- $(\mathcal{L}^p)^* = \mathcal{L}^q$ where $\frac{1}{p} + \frac{1}{q} = 1$
- $(\ell^1)^* = \ell^\infty$, but $(\ell^\infty)^* \supset \ell^1$, not equal

Section 3

Optimization in vector spaces

Theorem (First order condition for maximization with equality constraints)

Let $f : U \to \mathbb{R}$ and $h : U \to W$ be continuously differentiable on $U \subseteq V$, where V and W are Banach spaces. Suppose $x^* \in \mathrm{interior}(U)$ is a local maximizer of f on U subject to $h(x) = 0$. Suppose that $Dh_{x^*} : V \to W$ is onto. Then, there exists $\mu^* \in W^*$ such that for

$$L(x, \mu) = f(x) - \mu h(x).$$

we have

$$\begin{aligned} D_x L(x^*, \mu^*) =& Df_{x*} - \mu^* Dh_{x^*} = 0 \\ D_\mu L(x^*, \mu^*) =& h(x^*) = 0_W \end{aligned}$$

Section 4

Optimal control

Optimal control

Classic continuous time optimal control problem:

$$\max_{x(t),y(t)} \int_0^T F(x(t),y(t),t)dt$$

s.t.

$$\frac{dy}{dt} = g(x(t),y(t),t) \forall t \in [0,T]$$

$$y(0) = y_0$$

- x is called a control variable
- y is called a state variable
- V = space of pairs of functions $(x(t),y(t))$ from $[0,T]$ to $\mathbb{R}$
- Objective $f(x,y) = \int_0^T F(x(t),y(t),t)dt$, $f : V \to \mathbb{R}$
- Constraint $h(x,y) : V \to W$ = functions from $[0,T]$ to $\mathbb{R}$

$$h(x,y)(t) = \frac{dy}{dt}(t) - g(x(t),y(t),t)$$

- Multipliers $\mu \in W^*$, $\mu_0 \in \mathbb{R}^* = \mathbb{R}$

$$\max_{x(t),y(t)} \int_0^T F(x(t), y(t), t)dt$$
$$\text{s.t.} \frac{dy}{dt} = g(x(t), y(t), t) \forall t \in [0, T]$$
$$y(0) = y_0$$

- Lagrangian:

$$L(x, y, \mu, \mu_0) = \int_0^T F(x(t), y(t), t)dt - \mu\left(\frac{dy}{dt} - g(x, y, \cdot)\right) -$$
$$- \mu_0(y(0) - y_0)$$

- Assume that $\mu(w) = \int_0^T w(t)\lambda(t)dt$
- First order conditions are

$$[x]: \quad D_x f_{x^*,y^*} - \mu D_x h_{x^*,y^*} = 0 \tag{3}$$
$$[y]: \quad D_y f_{x^*,y^*} - \mu D_y h_{x^*,y^*} - \mu_0 = 0 \tag{4}$$
$$[\mu]: \quad h(x^*, y*) = 0 \tag{5}$$

- Rewrite first order conditions as

$$
\begin{aligned}
[x]:\ 0 =& \int_0^T v(t)\left(\frac{\partial F}{\partial x}(x(t),y(t),t) + \frac{\partial g}{\partial x}(x(t),y(t),t)\lambda(t)\right)dt \\
[y]:\ 0 =& \int_0^T v(t)\left(\frac{\partial F}{\partial y} + \frac{\partial g}{\partial y}\lambda(t)\right)dt + \\
& + \int_0^T \frac{d\lambda}{dt}(t)v(t)dt - \lambda(T)v(T) + \lambda(0)v(0) - \mu_0 v(0) \\
[\mu]:\ \frac{dy}{dt} =& g(x(t),y(t),t)
\end{aligned}
$$

- Equivalently, $\forall t \in [0, T]$

$$
\begin{aligned}
[x]: && 0 =& \frac{\partial F}{\partial x}(x(t),y(t),t) + \frac{\partial g}{\partial x}(x(t),y(t),t)\lambda(t) \\
[y]: && -\frac{d\lambda}{dt}(t) =& \frac{\partial F}{\partial y}(x(t),y(t),t) + \frac{\partial g}{\partial y}(x(t),y(t),t)v(t)\lambda(t) \\
[\mu]: && \frac{dy}{dt} =& g(x(t),y(t),t)
\end{aligned}
$$

and $\lambda(T) = 0$ and $\mu_0 = \lambda(0)$

Theorem (Pontryagin's maximum principle)

Consider

$$\max_{x,y \in U \subseteq X \times Y} \int_0^T F(x(t), y(t), t)dt$$

$$s.t. \tag{6}$$

$$\frac{dy}{dt} = g(x(t), y(t), t) \forall t \in [0, T] \tag{7}$$

$$y(0) = y_0.$$

where X and Y are some Banach spaces of differentiable functions from $[0, T]$ to $\mathbb{R}$, and $F, g : \mathbb{R}^2 \times [0, T] \to \mathbb{R}$ are continuously differentiable. Define the Hamiltonian as

$$H(x, y, \lambda, t) = F(x(t), y(t), t) + \lambda(t) g(x(t), y(t), t).$$

If x^* and y^* are a local constrained maximum of (7) in the interior of U, then there exists $\lambda^*(t)$ such that

$$[x] : \quad 0 = \frac{\partial H}{\partial x}(x^*, y^*, \lambda^*, t)$$

Application: optimal contracting with a continuum of types

$$\max_{q(\theta),T(\theta)} \int_{\theta_l}^{\theta_h} \left[T(\theta) - cq(\theta)\right] f_\theta(\theta) d\theta$$

s.t.

$$\theta\nu\left(q(\theta)\right) - T(\theta) \geq 0 \forall\theta \tag{8}$$

$$\theta\nu\left(q(\theta)\right) - T(\theta) \geq \max_{\tilde{\theta}} \theta\nu\left(q(\tilde{\theta})\right) - T(\tilde{\theta}) \forall\theta \tag{9}$$

- Can show equivalent to

$$\max_{q(\theta),T(\theta)} \int_{\theta_l}^{\theta_h} \left[T(\theta) - cq(\theta)\right] f_\theta(\theta) d\theta$$

s.t.

$$\theta_l \nu\left(q(\theta_l)\right) - T(\theta_l) \geq 0 \tag{10}$$

$$\theta \nu'(q(\theta))q'(\theta) - T'(\theta) = 0 \tag{11}$$

$$\frac{dq(\theta)}{d\theta} \geq 0 \tag{12}$$

- Manipulate first order conditions to eventually show

$$\left(\theta - \frac{1 - F_\theta(\theta)}{f_\theta(\theta)}\right) \nu'(q(\theta)) = c$$

Discrete time optimal control

$$\max_{x_t, y_y} \sum_{t=0}^{\infty} F(x_t, y_t, t)$$

s.t.

$$y_{t+1} - y_t = g(x_t, y_t, t)$$

First order conditions can be written as:

$$[x]: \quad \frac{\partial F}{\partial x}(x_t, y_t, t) + \mu_t \frac{\partial g}{\partial x}(x_t, y_t, t) = 0$$

$$[y]: \quad \mu_{t+1} - \mu_t = \frac{\partial F}{\partial y}(x_t, y_t, t) + \mu_t \frac{\partial g}{\partial y}(x_t, y_t, t) = 0$$

$$[\mu]: \quad y_{t+1} - y_t = g(x_t, y_t, t)$$

and $y_0 = 0$

- Same as continuous time, but with $\mu_{t+1} - \mu_t$ instead of $\frac{d\lambda}{dt}$ and $y_{t+1} - y_t$ instead of $\frac{dy}{dt}$

Section 5

Dynamic programming

Dynamic programming

- Optimal control: characterize maximizer x^*
- Dynamic programming: characterize maximal function value

Finite horizon

- Example: finite horizon consumption and savings

$$\max_{c_t, s_t} \sum_{t=0}^{T} \beta^t u(c_t) \text{ s.t. } s_{t+1} = (1 + r_t)(s_t - c_t)$$

  s_0 given and $s_{T+1} = 0$
- Backward induction:
  - Time T

$$V_T(s_T) \equiv \max_{c_T} u(c_T) \text{ s.t. } s_{T+1} = (1 + r_T)(s_T - c_T) = 0$$

  - $T - 1$

$$V_{T-1}(s_{T-1}) \equiv \max_{c_{T-1}, s'} u(c_{T-1}) + \beta V_T(s')$$
$$\text{s.t.} s' = (1 + r_{T-1})(c_{T-1} - s_{T-1})$$

  - In general, Bellman equation

$$V_t(s) = \max_{c_t, s'} u(c_t) + \beta V_{t+1}(s') \text{ s.t.} s' = (1 + r_t)(c_t - s) \tag{13}$$

Principle of optimality: "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." (Bellman and Dreyfus 1962)

Example: exercising a stock option

- Option contract: any time in next T periods, can buy a share at price p
- Market price at time t is x_t, $x_t = x_{t-1} + \epsilon_t$, ϵ_t i.i.d. with $\mathrm{E}|\epsilon_t|$ finite
- Problem: choose selling time s to maximize expected $x_s - p$
- Backward induction:
    - $V_T(x_T) = \max\{x_T - p, 0\}$
    - $V_{T-1}(x_{T-1}) = \max\{x_{T-1} - p, \mathrm{E}[V_T(x_T)|x_{T-1}]\} = \max\{x_{T-1} - p, \mathrm{E}[V_T(x_{T-1} + \epsilon_{T-1})]\}$
    - $V_t(x) = \max\{x - p, \mathrm{E}[V_{t+1}(x + \epsilon)]\}$

Example: exercising a stock option

- Use induction to show:
    1. If $t < \tau$ then $V_t(x) \geq V_\tau(x)$
    2. If $x < y$ then $V_t(x) - x \geq V_t(y) - y$
    3. $V_t(x)$ is continuous
- 2, 3, and $F(x) \geq x - p$ imply should use option at time t if $x \geq a_t$
- 1 implies $a_1 \geq a_2 \geq ... \geq a_T$

Infinite horizon

- Example:

$$\max_{c_t, s_t} \sum_{t=0}^{\infty} \beta^t u(c_t) \text{ s.t. } s_{t+1} = (1 + r_t)(s_t - c_t)$$

- Cannot use backward induction
- If problem does not depend on t,

$$V(s) = \max_{c, s'} u(c) + \beta V(s') \text{ s.t. } s' = (1 + r)(s - c)$$

- When will V not depend on t? Or will $\exists V$ that solves the above equation? And when does that V match the original problem?

Existence of value function

- Example:

$$\max_{c_t, s_t} \sum_{t=0}^{\infty} \beta^t u(c_t, s_t)$$
$$s_{t+1} = g(c_t, s_t),$$

where $c \in \mathbb{R}$, $s \in \mathbb{R}$, $0 < \beta < 1$, and $u, g : \mathbb{R}^2 \to \mathbb{R}$.

- Finding value function:
    - Guess $v_0 : \mathbb{R} \to \mathbb{R}$
    - Set

$$v_1(s) = \max_{c,s'} u(c, s) + \beta v_0(s') \text{ s.t. } s' = g(c, s)$$

    - Repeat: $v_j(s) = T(v_{j-1})(s)$ with

$$T(v)(s) = \max_{c,s'} u(c, s) + \beta v(s') \text{ s.t. } s' = g(c, s)$$

    - Recall: contraction mappings have unique fixed points
    - If T is a contraction, then V exists and $v_j \to V$

Existence of value function

- Want to show T is contraction mapping, i.e.

$$\left\| T(v) - T(v') \right\| \leq c \left\| v - v' \right\|$$

for $c < 1$

- Sufficient conditions:
  - Define space for v: u bounded, $u(c, s) \leq M$, then $v(s) \leq \frac{M}{1-\beta}$, i.e. $v \in \mathcal{L}^\infty(\mathbb{R})$, and $\|v\| = \sup_{x \in \mathbb{R}} |v(s)|$
  - Assure solution to

$$\max_{c_i, s_i'} u(c_i, s) + \beta v_i(s_i') \text{ s.t. } s_i' = g(c_i, s).$$

exists: assume c, s' in compact set

- Other more general conditions known

Existence of value function

- Show T contraction:
  - Note:

$$
\begin{aligned}
T(v_0)(s) - T(v_1)(s) = & (u(c_0, s) + \beta v_0(s_0')) - \\
& - (u(c_1, s) + \beta v_1(s_1'))
\end{aligned}
$$

    where c_i, s_i' are maximizer to Bellman equation with v_i
  - By maximization:

$$
Tv_0(s) = u(c_0, s) + \beta v_0(s_0') \geq u(c_1, s) + \beta v_0(s_1')
$$

    and vice versa
  - Therefore,

$$
T(v_0)(s) - T(v_1)(s) \geq \beta(v_0(s_1') - v_1(s_1'))
$$

    and

$$
T(v_0)(s) - T(v_1)(s) \leq \beta(v_0(s_0') - v_1(s_0'))
$$

  - Conclude:

$$
\begin{aligned}
\| T(v_0) - T(v_1) \| = & \sup_s |T(v_0)(s) - T(v_1)(s)| \\
\leq & \sup_s |\beta(v_0(s) - v_1(s))| = \beta \| v_0 - v_1 \|
\end{aligned}
$$

Example: linear-quadratic

$$\min_{x_t,u_t} \sum_{t=0}^{\infty} \beta^t (x_t - a)^2 \text{ s.t. } x_{t+1} = bx_t + cu_t$$

- Strategy: guess and check functional form of $V(x)$
  - Does not always work, but when it does V usually has same form as the objective function
  - Many problems have no closed form expression for V
- Guess $V(x) = v_0 + v_1 x + v_2 x^2$
- Solve $\min_u (x-a)^2 + \beta V(bx + cu)$, get $u^*(x) = \frac{v_1}{v_2} - \frac{b}{c}x$
- Substitute:

$$\begin{aligned} V(x) =& (x-a)^2 + \beta V(bx + cu^*(x)) \\ v_0 + v_1 x + v_2 x^2 =& (x-a)^2 + \beta \left[v_0 + \frac{cv_1^2}{v_2} + \frac{v_1^2 c^2}{v_2} \right] \end{aligned}$$

solve for v_0, v_1, v_2 (will not be possible if guess incorrect)

Example: optimal growth

- log preferences, Cobb-Douglas production

$$\max_{\{c_t\}_{t=0}^\infty} \sum_{t=0}^\infty \delta^t \log(c_t)$$
$$\text{s.t. } c_t + k_{t+1} = k_t^\alpha.$$

- Bellman equation:

$$v(k) = \max_{k' \in [0, k^\alpha]} \log(k^\alpha - k') + \delta v(k')$$

- Could guess and verify $v(k) = c_0 + c_1 \log(k)$
- If cannot guess, can try iterating the Bellman operator

Example: optimal growth I

- Set $v_0(k) = 0$ (or anything)
- Make sequence $v_i = T(v_{i-1})$
- v_1:

$$\begin{aligned} v_1(k) = & \max_{k' \in [0, k^\alpha]} \log(k^\alpha - k') + \delta v_0(k) \\ = & \max_{k' \in [0, k^\alpha]} \log(k^\alpha - k') \\ = & \alpha \log k. \end{aligned}$$

- v_2:

$$\begin{aligned} v_2(k) = & \max_{k' \in [0, k^\alpha]} \log(k^\alpha - k') + \delta v_1(k) \\ = & \max_{k' \in [0, k^\alpha]} \log(k^\alpha - k') + \delta \alpha \log(k') \\ = & c_2 + (\alpha + \delta \alpha^2) \log k, \end{aligned}$$

Example: optimal growth II

- Etc.

$$
\begin{aligned}
v_3(k) =& c_3 + (\alpha + \delta\alpha^2 + \delta^2\alpha^3)\log k \\
v_4(k) =& c_4 + (\alpha + \delta\alpha^2 + \delta^2\alpha^3 + \delta^3\alpha^4)\log k \\
& \vdots \\
v_i(k) =& c_i + \alpha \sum_{j=0}^{i} (\alpha\delta)^j \log(k)
\end{aligned}
$$

- Take, limit $i \to \infty$,

$$
v(k) = C + \frac{\alpha}{1-\alpha\delta}\log k
$$