# 16. Pontryagin's maximum principle

Recall, for $(s,x) \in \tilde{S}$, we seek to

minimize

$$V^u(s,x) = \int_s^\tau c(t, x_t, u_t)\,dt + C(\tau, x_\tau)$$

subject to

$$\dot{x}_t = b(t, x_t, u_t), \quad x_s = x, \quad (\tau, x_\tau) \in \tilde{D}$$

where

$$\tau = \inf\{t \geq 0 : x_t \in D\}.$$

We make some regularity assumptions

- the stopping set $D$ is a hyperplane,

  thus $D = \{y\} + \Sigma$ for some $y \in \mathbb{R}^d$ and some vector subspace $\Sigma$ of $\mathbb{R}^d$.

- $b : \tilde{S} \times A \to \mathbb{R}^d$, $c : \tilde{S} \times A \to \mathbb{R}$, $C : \tilde{D} \to \mathbb{R}$

  are differentiable on $\tilde{S}$ and $\tilde{D}$ with continuous derivatives

Recall, either $\tilde{S} = [0,T] \times \mathbb{R}^d$, $\tilde{D} = \{T\} \times D$ (fixed horizon)

or $\tilde{S} = \mathbb{R}^+ \times S$, $\tilde{D} = \mathbb{R}^+ \times D$, $D = \partial S$

(unconstrained horizon).

Pontryagin's maximum principle states that, if $(x_t, u_t)_{t\le\tau}$ is optimal, then there exist adjoint paths $(\lambda_t)_{t\le\tau}$ in $\mathbb{R}^d$ and $(\mu_t)_{t\le\tau}$ in $\mathbb{R}$ with the following properties: for all $t\le\tau$,

(i) $\lambda_t^T b(t, x_t, \cdot) - c(t, x_t, \cdot) + \mu_t \le 0$, with equality at $u_t$;

(ii) $\dot{\lambda}_t^T = -\lambda_t^T \nabla b(t, x_t, u_t) + \nabla c(t, x_t, u_t)$;

(iii) $\dot{\mu}_t = -\lambda_t^T \dot{b}(t, x_t, u_t) + \dot{c}(t, x_t, u_t)$;

(iv) $\dot{x}_t = b(t, x_t, u_t)$;

moreover the following transversality conditions hold:

(v) $(\lambda_\tau^T + \nabla C(\tau, x_\tau))\sigma = 0$ for all $\sigma \in \Sigma$;

(vi) $\mu_\tau + \dot{C}(\tau, x_\tau) = 0$ (unconstrained horizon case only).

- A good scheme to remember (i) to (iv) is to define the <u>Hamiltonian</u>

$$H(t, x, u, \lambda) = \lambda^T b(t, x, u) - c(t, x, u).$$

Then we have

(i) $0 = \frac{\partial H}{\partial u}$, (ii) $\dot{\lambda} = -\frac{\partial H}{\partial x}$, (iii) $\dot{\mu} = -\frac{\partial H}{\partial t}$, (iv) $\dot{x} = \frac{\partial H}{\partial \lambda}$

(may fail when $A$ is an interval and $H$ is maximized at an endpoint

- In the unconstrained horizon case, if $b$, $c$ and $C$ are time-independent, then (vi) and (iii) imply $\mu_t \equiv 0$.

## Example - bringing a particle to rest at $O$ in minimal time

A particle moves on a line, with initial position $q_0$ and velocity $p_0$. We can apply a force which imparts to the particle an acceleration $a \in [-1, 1]$.
How can we bring the particle to rest at $O$ in the shortest time?

## Example – Miss Prout

Miss P. holds the entire remaining stock of Cambridge elderberry wine for the vintage year 1959.

If she releases it at rate $u$, then she realizes a unit price

$$p(u) = \max\{0, 1 - u/2\}.$$

She holds an amount $x$ at time $0$.

What is her maximal total discounted return

$$\int_0^\infty e^{-\alpha t} u_t \, p(u_t) \, dt$$

and how should she achieve it?

## Example – insect optimization

Workers $w_t$

Queens $q_t$

If workers devote a proportion $u_t$ of their effort to producing more workers, then

$$\dot{w}_t = a u_t w_t - b w_t \,, \quad \dot{q}_t = (1-u_t) w_t \,.$$

Here $a, b \in (0, \infty)$ with $a > b$.

How is the number of queens maximized by time $T$?

Recall the method of Lagrange multipliers:

To solve maximize $f(x)$,

subject to $g(x) = b \in \mathbb{R}^d$,

- define the Lagrangian

$$L(x, \lambda) = f(x) - \lambda^T (g(x) - b), \quad \lambda \in \mathbb{R}^d,$$

- find $x(\lambda)$ to maximize $L(x, \lambda)$ for each $\lambda$,
- seek $\lambda^*$ so that $g(x(\lambda^*)) = b$,

then $x(\lambda^*)$ is the desired maximizer.

We now make a non-rigorous analogue of this to derive PMP, for simplicity in the time-homogeneous case with fixed time horizon $T$.

We wish to maximize $\displaystyle -\int_0^T c(x_t, u_t)\,dt - C(x_T)$

subject to $\dot{x}_t = b(x_t, u_t)$, $0 \le t \le T$.

Define for a path $(\lambda_t)_{0 \le t \le T}$ in $\mathbb{R}^d$

$$L(x, \lambda) = \int_0^T \{-c(x_t, u_t) - \lambda_t^T(\dot{x}_t - b(x_t, u_t))\}\,dt - C(x_T)$$

$$= -\lambda_T^T x_T + \lambda_0^T x_0 + \int_0^T \{\dot{\lambda}_t^T x_t + \lambda_t^T b(x_t, u_t) - c(x_t, u_t)\}\,dt - C(x_T)$$

by an integration by parts.

$$L(x,\lambda) = -\lambda_T^T x_T + \lambda_0^T x_0 + \int_0^T \{ \dot{\lambda}_t^T x_t + \lambda_t^T b(x_t, u_t) - c(x_t, u_t) \} dt - C(x_T)$$

Fix $\lambda$. At a maximum in $x$, we might expect

$$0 = \frac{\partial L}{\partial x_t} = \dot{\lambda}_t^T + \lambda_t^T \nabla b(x_t, u_t) - \nabla c(x_t, u_t)$$

which is the adjoint equation

and

$$0 = \frac{\partial L}{\partial x_T} \cdot \sigma = (-\lambda_T^T - \nabla C(X_T)) \cdot \sigma \quad , \quad \sigma \in \Sigma$$

which is the tranversality condition.

- We will now give a proof of PMP assuming the existence of a twice continuously differentiable solution of the HJB equation.

  In many cases of interest this assumption is false, but at least the proof will provide some rigorous basis for using PMP

- We assume in the next result that the action space $A$ is an open set in $\mathbb{R}^p$ and that $c, b : \tilde{S} \times A \to \mathbb{R}$ and $C : \tilde{D} \times A \to \mathbb{R}$ are continuously differentiable.

Proposition 16.1

Suppose that there exist

- $F : \tilde{S} \cup \tilde{D} \to \mathbb{R}$ twice continuously differentiable
- $u : \tilde{S} \to A$ continuously differentiable

such that $F = C$ on $\tilde{D}$ and, for all $(t,x) \in \tilde{S}$,

$$c(t,x,a) + \dot{F}(t,x) + \nabla F(t,x) b(t,x,a) \geq 0, \quad a \in A$$

with equality at $a = u(t,x)$.

Fix $(0,x) \in \tilde{S}$ and define $(x_t, u_t)_{t \leq \tau}$ by

$$\dot{x}_t = b(t, x_t, u(t,x_t)), \quad x_0 = x, \quad u_t = u(t,x_t), \quad \tau = \inf\{t \geq 0 : (t,x_t) \in \tilde{D}\}$$

Assume that $\tau < \infty$. Set $\mu_t = -\dot{F}(t,x_t)$, $\lambda_t^T = -\nabla F(t,x_t)$, then

$$\dot{\lambda}_t^T = -\lambda_t^T \nabla b(t,x_t,u_t) + \nabla c(t,x_t,u_t), \quad \mu_t = -\lambda_t^T \dot{b}(t,x_t,u_t) + \dot{c}(t,x_t,u_t)$$

and, for any $\sigma \in \Sigma$, $(\lambda_\tau^T + \nabla C)(\tau, x_\tau).\sigma = 0$

and, in the time-unconstrained case, $\mu_\tau + \dot{C}(\tau, x_\tau) = 0$.

Proof Consider

$$J(t,x,a) = c(t,x,a) + \dot F(t,x) + \nabla F(t,x)\, b(t,x,a), \quad (t,x) \in \tilde S,\ a \in A.$$

Then $J$ is continuously differentiable, $J \geq 0$ and $J(t,x,u(t,x)) = 0$ for all $(t,x) \in \tilde S$. Since $A$ is open, $\left(\frac{\partial J}{\partial a}\right)(t,x,u(t,x)) = 0$, and so

$$\nabla J(t,x,u(t,x)) = \frac{\partial}{\partial x}\big(J(t,x,u(t,x))\big) = 0, \quad \dot J(t,x,u(t,x)) = \frac{\partial}{\partial t}\big(J(t,x,u(t,x))\big) = 0,$$

So, for $a = u(t,x)$ we have

$$\nabla c(t,x,a) + \nabla F(t,x) \nabla b(t,x,a) + \{\nabla \dot F(t,x) + \nabla^2 F(t,x)\, b(t,x,a)\} = 0,$$

$$\dot c(t,x,a) + \nabla F(t,x)\, \dot b(t,x,a) + \{\ddot F(t,x) + \nabla \dot F(t,x)\, b(t,x,a)\} = 0.$$

Hence

$$\dot\lambda_t^T = -\nabla \dot F(t,x_t) - \nabla^2 F(t,x_t)\,\dot x_t = \nabla c(t,x_t,u_t) + \overset{(-\lambda_t^T)}{\cancel{\nabla F(t,x_t)}}\, \nabla b(t,x_t,u_t),$$

$$\dot\mu_t = -\ddot F(t,x_t) - \nabla \dot F(t,x_t)\,\dot x_t = \dot c(t,x_t,u_t) + \nabla F\cancel{(t,x_t)}\, \dot b(t,x_t,u_t).$$

Finally, we can differentiate the equality $V = C$ at $(\tau, x_\tau)$ in $x$ in the direction $\sigma$ and, in the unconstrained case, in $t$. $\square$ 16/4