Optimal Control

Harald Lang

Continuous case

Consider the problem of finding functions $x(t) = (x_1(t), \ldots, x_n(t))$ and $u(t) = (u_1(t), \ldots, u_m(t))$ so as to maximize the integral

$$\int_a^b f(x(t), u(t), t)\, dt$$

subject to

$$\dot{x}_j(t) = g_j(x(t), u(t), t) \quad j = 1, \ldots, n$$

$$\int_a^b k_j(x(t), u(t), t)\, dt = Q_j \quad j = 1, \ldots, J_1$$

$$h_j(x(t), u(t), t) = 0 \quad j = 1, \ldots, J_2$$

$$u(t) \in \Gamma_t \quad \Gamma_t \text{ some convex subset of } \mathbf{R}^m$$

Of course, $J_1 = 0$ means that there are no integral restrictions, and similarly if $J_2 = 0$. Define the *Hamiltonian* H as

$$\begin{aligned} H(x(t), u(t), \lambda(t), \mu, \nu(t), t) = & f(x(t), u(t), t) + \lambda(t) \cdot g(x(t), u(t), t) \\ & + \mu \cdot k(x(t), u(t), t) + \nu(t) \cdot h(x(t), u(t), t) \end{aligned}$$

where we of course have used vector notation: $g(\cdot) = (g_1(\cdot), \ldots, g_n(\cdot))$; $\lambda(t) = (\lambda_1(t), \ldots, \lambda_n(t))$, etc. Note that the *Lagrangean multipliers* λ_j and ν_j are functions of t, whereas the μ_j's are constants.

Assumption: *The Hamiltonian H is concave in the variables (x, u) and continuously differentiable wrt x.*

Lemma 1. *Let $u \in \operatorname*{argmax}_{u \in \Gamma_t} H(x, u, \lambda, \mu, \nu, t)$. Then, for any $\tilde{u} \in \Gamma_t$,*

$$H(x, u, \lambda, \mu, \nu, t) - H(\tilde{x}, \tilde{u}, \lambda, \mu, \nu, t) \geq H_x(x, u, \lambda, \mu, \nu, t) \cdot [x - \tilde{x}]$$

Proof. Let $\Delta x \equiv \tilde{x} - x, \Delta u \equiv \tilde{u} - u$. For any positive integer m,

$$\begin{aligned} &H(x + \frac{\Delta x}{m}, u + \frac{\Delta u}{m}, \lambda, \mu, \nu, t) \\ &\qquad \geq H(x, u, \lambda, \mu, \nu, t) + \frac{1}{m}[H(\tilde{x}, \tilde{u}, \lambda, \mu, \nu, t) - H(x, u, \lambda, \mu, \nu, t)] \\ &\qquad \geq H(x, u + \frac{\Delta u}{m}, \lambda, \mu, \nu, t) + \frac{1}{m}[H(\tilde{x}, \tilde{u}, \lambda, \mu, \nu, t) - H(x, u, \lambda, \mu, \nu, t)] \end{aligned}$$

The first inequality comes from the concavity, the second from the definition of u; note that $u + \frac{\Delta u}{m} \in \Gamma_t$ since Γ_t is convex. Hence

$$\begin{aligned} &H(x, u, \lambda, \mu, \nu, t) - H(\tilde{x}, \tilde{u}, \lambda, \mu, \nu, t) \\ &\qquad \geq m[H(x, u + \frac{\Delta u}{m}, \lambda, \mu, \nu, t) - H(x + \frac{\Delta x}{m}, u + \frac{\Delta u}{m}, \lambda, \mu, \nu, t)] \\ &\qquad = -H_x(\widehat{x}_m, u + \frac{\Delta u}{m}, \lambda, \mu, \nu, t) \cdot \Delta x \end{aligned}$$

where $\widehat{x}_m$ is some point on the straight line connecting x and $x + \frac{\Delta x}{m}$; the equality comes from the mean value theorem of calculus. Now, letting $m \to \infty$ using the fact that H_x is continuous, we get the result of the lemma; *Q.E.D.*

It is easy to verify that

$$\int_a^b f(x(t), u(t), t)\, dt = \int_a^b [H(\cdot) - \lambda(t) \cdot \dot{x}(t)]\, dt - \mu Q$$

Hence, if $(x(t), u(t))$ is a fixed feasible pair with

$$u(t) \in \operatorname*{argmax}_{u \in \Gamma_t} H(x(t), u(t), \lambda(t), \mu, \nu, t)$$

and $(\tilde{x}(t), \tilde{u}(t))$ is any feasible pair, then

$$\begin{aligned} &\int_a^b f(x(t), u(t), t)\, dt - \int_a^b f(\tilde{x}(t), \tilde{u}(t), t)\, dt \\ &= \int_a^b \{H(x(t), u(t), \lambda(t), \mu, \nu(t), t) - H(\tilde{x}(t), \tilde{u}(t), \lambda(t), \mu, \nu(t), t)\}\, dt \\ &\quad - \int_a^b \lambda(t) \cdot [\dot{x}(t) - \dot{\tilde{x}}(t)]\, dt \\ &\geq \int_a^b H_x(x(t), u(t), \lambda(t), \mu, \nu(t), t) \cdot [x(t) - \tilde{x}(t)]\, dt - \int_a^b \lambda(t) \cdot [\dot{x}(t) - \dot{\tilde{x}}(t)]\, dt \\ &= \int_a^b [H_x(x(t), u(t), \lambda(t), \mu, \nu(t), t) + \dot{\lambda}(t)] \cdot [x(t) - \tilde{x}(t)]\, dt \\ &\quad - \big[\lambda(t) \cdot (x(t) - \tilde{x}(t))\big]_a^b \end{aligned}$$

The following conditions on $(x(t), u(t), \lambda(t), \mu, \nu(t))$ are obviously sufficient for the last expression in the chain of relations above to be 0, and hence *sufficient conditions for $(x(t), u(t))$ to be a solution to the maximization problem:*

(Ex$_j$) $\quad H_{x_j}(x(t), u(t), \lambda(t), \mu, \nu(t), t) + \dot{\lambda}_j(t) = 0 \quad j = 1, \ldots, n$

(Mu) $\quad u(t) \in \operatorname*{argmax}_{u \in \Gamma_t} H(x(t), u(t), \lambda(t), \mu, \nu(t), t) \quad a \leq t \leq b$

for any $j = 1, \ldots, n$, either $x_j(b) = x_j^b$ is given, or else

(TUx$_j$) $\lambda_j(b) = 0$

for any $j = 1, \ldots, n$, either $x_j(a) = x_j^a$ is given, or else

(TLx$_j$) $\lambda_j(a) = 0$

If there is no restriction $x_j(b) = x_j^b$, then b is called a *free boundary* for x_j, and the condition *(TUx$_j$)* and *(TLx$_j$)* are called *transversality conditions.* The E-equations are called *Euler equations.*

Infinite horizon

Let in the previous case $b = \infty$. The only impact this has on the analysis above is on the last expression $-[\lambda(t) \cdot (x(t) - \tilde{x}(t))]_a^b$. A sufficient condition for this expression to be ≥ 0 as $b \to \infty$ is, in addition to *(TLx$_j$)* above

(TU$^\infty$x$_j$) $\liminf_{t \to \infty} \lambda(t) \cdot \tilde{x}(t) \geq 0$ *for all feasible $\tilde{x}(t)$ and* $\lim_{t \to \infty} \lambda(t) \cdot x(t) = 0$

The envelope theorem

Let $V = \max \int_a^b f(x(t), u(t), t)\, dt$ and let α be a parameter which does not influence Q, a, b, Γ_t, x^a, if a is not a free boundary, x^b if b is not a free boundary. Assume that the sufficient conditions given above are satisfied and that u is uniquely defined by *(Mu)* and locally bounded as a function of α. Then

$$\frac{dV}{d\alpha} = \int_a^b H_\alpha(x(t), u(t), \lambda(t), \mu, \nu(t), t)\, dt$$

$$\frac{dV}{da} = -H(x(a), u(a), \lambda(a), \mu, \nu(a), a)$$

$$\frac{dV}{db} = H(x(b), u(b), \lambda(b), \mu, \nu(b), b)$$

$$\frac{dV}{dx_j^a} = \lambda_j(a) \ \ \textit{(if a is not a free boundary)}$$

$$\frac{dV}{dx_j^b} = -\lambda_j(b) \ \ \textit{(if b is not a free boundary)}$$

$$\frac{dV}{dQ_j} = -\mu_j \ \ j = 1, \ldots, J_1$$

Lemma 2. *If β is any parameter (i.e., β may influence a, x_j^a, Q, etc.), then*

$$\int_a^b [H_x \cdot x_\beta + H_\lambda \cdot \lambda_\beta + H_\mu \cdot \mu_\beta + H_\nu \cdot \nu_\beta - \lambda_\beta \cdot \dot{x} - \lambda \cdot \dot{x}_\beta] \, dt - \mu_\beta \cdot Q$$
$$= \lambda(a) \cdot x_\beta(a) - \lambda(b) \cdot x_\beta(b)$$

Proof. The term $H_\mu = k(x, u, t)$ whose integral is Q and $H_\nu = h(x, u, t)$ which is $\equiv 0$. Hence we are left with

$$\int_a^b [H_x \cdot x_\beta + H_\lambda \cdot \lambda_\beta - \lambda_\beta \cdot \dot{x} - \lambda \cdot \dot{x}_\beta] \, dt = \int_a^b [(H_x + \dot{\lambda}) \cdot x_\beta + (H_\lambda - \dot{x}) \cdot \lambda_\beta] \, dt - [\lambda(t) \cdot x_\beta(t)]_a^b$$

The integral is $= 0$, since $H_x + \dot{\lambda} = 0$ by *(Ex)* and $H_\lambda - \dot{x} = g(x, u, t) - \dot{x} = 0$. So we are left with $\lambda(a) \cdot x_\beta(a) - \lambda(b) \cdot x_\beta(b)$ *Q.E.D.*

Proof of the the envelope theorem.

By the "usual" envelope theorem, we can treat u as constant when we differentiate, hence

$$\frac{dV}{d\alpha} = \int_a^b H_\alpha \, dt + \int_a^b [H_x \cdot x_\alpha + H_\lambda \cdot \lambda_\alpha + H_\mu \cdot \mu_\alpha + H_\nu \cdot \nu_\alpha - \lambda_\alpha \cdot \dot{x} - \lambda \cdot \dot{x}_\alpha] \, dt - \mu_\alpha \cdot Q$$

Using lemma 2, we get

$$\frac{dV}{d\alpha} = \int_a^b H_\alpha \, dt + \lambda(a) \cdot x_\alpha(a) - \lambda(b) \cdot x_\alpha(b)$$

If a is a free boundary, then $\lambda(a) = 0$; if $x(a) = x^a$ is given, then $x_\alpha(a) = 0$, since x^a is independent of α by assumption. In any case, the product $\lambda(a) \cdot x_\alpha(a) = 0$, and similarly $\lambda(b) \cdot x_\alpha(b) = 0$, which proves the $dV/d\alpha$ part.

$$\begin{aligned}
\frac{dV}{db} &= \frac{d}{db}\Big\{\int_a^b [H(\cdot) - \lambda(t) \cdot \dot{x}(t)] \, dt - \mu \cdot Q\Big\} \\
&= H(x(b), u(b), \lambda(b), \mu, \nu, b) - \lambda(b) \cdot \dot{x}(b) \\
&\quad + \int_a^b [H_x \cdot x_b + H_\lambda \cdot \lambda_b + H_\mu \cdot \mu_b + H_\nu \cdot \nu_b - \lambda_b \cdot \dot{x} - \lambda \cdot \dot{x}_b] \, dt - \mu_b \cdot Q
\end{aligned}$$

Using lemma 2, we get

$$\frac{dV}{db} = H(x(b), u(b), \lambda(b), \mu, \nu, b) - \lambda(b) \cdot \dot{x}(b) + \lambda(a) \cdot x_b(a) - \lambda(b) \cdot x_b(b)$$

Here the term $\lambda(a) \cdot x_b(a) = 0$, by the same argument as in the previous case. If b is a free boundary, then $\lambda(b) = 0$. In order to analyze the situation when b is not a free boundary, we introduce the temporary notation $x(t; b)$ for the optimal function $x(t)$ given that the upper limit of integration is b. In this case $x(b; b) = \widehat{x}$ where $\widehat{x}$ is a number independent of b. Hence, by differentiation w.r.t. b, $\dot{x}(b,b)+x_b(b;b)=0$. We see that in all cases the sum $-\lambda(b)\cdot\dot{x}(b)-\lambda(b)\cdot x_b(b)=0$, which proves the dV/db part. Of course, the dV/da part is proven similarly.

Using lemma 2, we have

$$\begin{aligned}\frac{dV}{dx_j^a} &= \int_a^b \Big[H_x\cdot\frac{\partial x}{\partial x_j^a}+H_\lambda\cdot\frac{\partial\lambda}{\partial x_j^a}+H_\mu\cdot\frac{\partial\mu}{\partial x_j^a}+H_\nu\cdot\frac{\partial\nu}{\partial x_j^a}-\frac{\partial\lambda}{\partial x_j^a}\cdot\dot{x}-\lambda\cdot\frac{\partial\dot{x}}{\partial x_j^a}\Big]\,dt\\ &\quad -Q\cdot\frac{\partial x}{\partial x_j^a}\\ &= \lambda(a)\cdot\frac{\partial x(a)}{\partial x_j^a}-\lambda(b)\cdot\frac{\partial x(b)}{\partial x_j^a}\end{aligned}$$

Here either $\lambda(b)$ is $=0$ (if b is a free boundary) or $\dfrac{\partial x(b)}{\partial x_j^a}=0$ (if b is not a free boundary) Hence $\dfrac{dV}{dx_j^a}=\lambda(a)\cdot\dfrac{\partial x(a)}{\partial x_j^a}=\lambda_j(a)$. The formula for $\dfrac{dV}{dx_j^b}$ is of course proven similarly.

Finally, employing lemma 2 once again, we have

$$\begin{aligned}\frac{dV}{dQ} &= \frac{d}{dQ}\Big\{\int_a^b[H(\cdot)-\lambda(t)\cdot\dot{x}(t)]\,dt-\mu\cdot Q\Big\}\\ &= \int_a^b H_Q\,dt+\int_a^b[H_x\cdot x_Q+H_\lambda\cdot\lambda_Q+H_\mu\cdot\mu_Q+H_\nu\cdot\nu_Q-\lambda_Q\cdot\dot{x}-\lambda\cdot\dot{x}_Q]\,dt\\ &\quad -Q\cdot\mu_Q-\mu\\ &= \lambda(a)\cdot x_Q(a)-\lambda(b)\cdot x_Q(b)-\mu=-\mu\end{aligned}$$

The argument for the last equality is the same as in previous cases. *Q.E.D.*

Discrete case

Consider the problem of finding sequences $x_t=(x_t^1,\ldots,x_t^n)$ and $u_t=(u_t^1,\ldots,u_t^m)$ so as to maximize the sum

$$\sum_s^T f(x_t,u_t,t)$$

subject to the constraints

$$\begin{aligned}&x_{t+1}^j=g^j(x_t,u_t,t),\quad t=s,\ldots,T\quad j=1,\ldots,n\\ &\sum_s^T k^j(x_t,u_t,t)=Q^j\quad j=1,\ldots,J_1\\ &h^j(x_t,u_t,t)=0\quad t=s,\ldots,T\quad j=1,\ldots,J_2\\ &u_t\in\Gamma_t\quad \Gamma_t\text{ some convex subset of }\mathbf{R}^m\end{aligned}$$

Of course, $J_1=0$ means that there are no summation restrictions, and similarly if $J_2=0$. Define the *Hamiltonian*

$$H(x_t,u_t,\lambda_{t+1},\mu,\nu_t,t)=f(x_t,u_t,t)+\lambda_{t+1}\cdot g(x_t,u_t,t)+\mu\cdot k(x_t,u_t,t)+\nu_t\cdot h(x_t,u_t,t)$$

where we of course have used vector notation: $g(\cdot) = (g^1(\cdot), \ldots, g^n(\cdot))$; $\lambda_t = (\lambda_t^1, \ldots, \lambda_t^n)$, etc. Note that the *Lagrangean multipliers* λ^j and ν^j are functions of t, whereas the μ^j's are constants. It is easy to verify that

$$\sum_s^T f(x_t, u_t, t) = \sum_s^T [H(\cdot) - \lambda_{t+1} \cdot x_{t+1}] - \mu Q$$

Assumption: *The Hamiltonian H is concave in the variables (x, u) and continuously differentiable wrt x.*

By lemma 1, if $u_t \in \operatorname{argmax}_{u \in \Gamma_t} H(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t)$ and $(\tilde{x}_t, \tilde{u}_t)$ is any feasible pair, then

$$H(x, u, \lambda, \mu, \nu, t) - H(\tilde{x}, \tilde{u}, \lambda, \mu, \nu, t) \geq H_x(x, u, \lambda, \mu, \nu, t) \cdot [x - \tilde{x}]$$

Hence, if (x_t, u_t) is a fixed feasible pair such that $u_t \in \operatorname{argmax}_{u \in \Gamma_t} H(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t)$ and $(\tilde{x}_t, \tilde{u}_t)$ is any feasible pair, then

$$\begin{aligned}
&\sum_s^T f(x_t, u_t, t) - \sum_s^T f(\tilde{x}_t, \tilde{u}_t, t) \\
&\qquad = \sum_s^T \{H(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) - H(\tilde{x}_t, \tilde{u}_t, \lambda_{t+1}, \mu, \nu_t, t)\} \\
&\qquad\quad - \sum_s^T \lambda_{t+1} \cdot [x_{t+1} - \tilde{x}_{t+1}] \\
&\qquad \geq \sum_s^T H_x(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) \cdot [x_t - \tilde{x}_t] - \sum_s^T \lambda_{t+1} \cdot [x_{t+1} - \tilde{x}_{t+1}] \\
&\qquad = \sum_s^T [H_x(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) - \lambda_t] \cdot [x_t - \tilde{x}_t] + \lambda_s \cdot [x_s - \tilde{x}_s] \\
&\qquad\quad - \lambda_{T+1} \cdot [x_{T+1} - \tilde{x}_{T+1}]
\end{aligned}$$

The following conditions on $(x_t, u_t, \lambda_t, \mu, \nu_t)$ are obviously sufficient for this expression to be $= 0$, and hence *sufficient conditions for (x_t, u_t) to be a solution to the maximization problem:*

(Ex^j) $\quad H_{x^j}(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) - \lambda_t^j = 0 \quad t = s, \ldots, T \quad j = 1, \ldots, n$

(Mu) $\quad u_t \in \operatorname{argmax}_{u \in \Gamma_t} H(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) \quad t = s, \ldots, T^*$

for any $j = 1, \ldots, n$, either x_j^{T+1} is given, or else

(TUx^j) $\quad \lambda_{T+1}^j = 0$

for any $j = 1, \ldots, n$, either x_s^j is given, or else

(TLx^j) $\quad \lambda_s^j = 0$

* *In contrast to the continuous case, this is not a necessary condition if H is not concave.*

Infinite horizon

Let in the previous case $T = \infty$. The only impact this has on the analysis above is on the last expression $-\lambda_{T+1} \cdot [x_{T+1} - \tilde{x}_{T+1}]$. A sufficient condition for this expression to be ≥ 0 as $T \to \infty$ is, in addition to *(TLx^j)* above,

$$(TU^\infty x^j) \quad \liminf_{t\to\infty} \lambda_t \cdot \tilde{x}_t \geq 0 \quad \text{for all feasible } \tilde{x}_t \text{ and } \quad \lim_{t\to\infty} \lambda_t \cdot x_t = 0$$

The envelope theorem

Let $V = \max \sum_s^T f(x_t, u_t, t)$ and let α be a parameter which does not influence Q, s, T, Γ_t or any of the boundary values of x, if any are given as constraints. Assume that the sufficient conditions given above are satisfied and that u_t^j is uniquely determined by *(Mu)* and locally bounded in α. Then

$$\frac{dV}{d\alpha} = \sum_s^T H_\alpha(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t)$$

$$\frac{dV}{dx_s^j} = \lambda_s^j \quad \text{if } x_s^j \text{ is given as a constraint}$$

$$\frac{dV}{dx_{T+1}^j} = -\lambda_{T+1}^j \quad \text{if } x_{T+1}^j \text{ is given as a constraint}$$

$$\frac{dV}{dQ_j} = -\mu^j \quad j = 1, \ldots, J_1$$

Proof. Using the "usual" envelope theorem as to the variation in u,

$$\begin{aligned}
\frac{dV}{d\alpha} &= \frac{d}{d\alpha} \sum_s^T f(x_t, u_t, t) \\
&= \frac{d}{d\alpha} \Big\{ \sum_s^T [H(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) - \lambda_{t+1} \cdot x_{t+1}] - \mu \cdot Q \Big\} \\
&= \sum_s^T H_\alpha(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) + \sum_s^T \Big\{ H_x \cdot \frac{dx_t}{d\alpha} + H_\lambda \cdot \frac{d\lambda_{t+1}}{d\alpha} + H_\mu \cdot \frac{d\mu}{d\alpha} \\
&\quad + H_\nu \cdot \frac{d\nu_t}{d\alpha} - \frac{d\lambda_{t+1}}{d\alpha} \cdot x_{t+1} - \lambda_{t+1} \cdot \frac{dx_{t+1}}{d\alpha} \Big\} - \frac{d\mu}{d\alpha} \cdot Q \\
&= \sum_s^T H_\alpha(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) + \sum_s^T [H_x - \lambda_t] \cdot \frac{dx_t}{d\alpha} + \sum_s^T [H_\lambda - x_{t+1}] \cdot \frac{d\lambda_{t+1}}{d\alpha} \\
&\quad + \Big\{ \Big(\sum_s^T H_\mu \Big) - Q \Big\} \cdot \frac{d\mu}{d\alpha} + \sum_s^T H_\nu \cdot \frac{d\nu_t}{d\alpha} + \lambda_s \cdot \frac{dx_s}{d\alpha} - \lambda_{T+1} \cdot \frac{dx_{T+1}}{d\alpha}
\end{aligned}$$

Using *(Ex)*, $H_\lambda - x_{t+1} = g(x_t, u_t, t) - x_{t+1} = 0$, $\sum_s^T H_\mu - Q = \sum_s^T k(x_t, u_t, t) - Q = 0$ and $H_\nu = h(\ldots) = 0$, we get

$$\frac{dV}{d\alpha} = \sum_s^T H_\alpha(x_t, u_t, \lambda_{t+1}, \mu, \nu_t, t) + \lambda_s \cdot \frac{dx_s}{d\alpha} - \lambda_{T+1} \cdot \frac{dx_{T+1}}{d\alpha}$$

Here, either $\lambda_s = 0$ (if s is a free boundary) or $dx_s/d\alpha = 0$ (if x_s is given); hence the λ_s-term $= 0$, and similarly the λ_{T+1}-term $= 0$. The other derivatives are shown similarly. *Q.E.D.*

Uncertainty

Let t be time, ξ_t a stochastic process, E_t the expectations operator conditional on I_t, the information set available at time t; in particular, $t' < t \Rightarrow I_{t'} \subseteq I_t$. Consider the problem of finding sequences $x_t = (x_t^1, \ldots, x_t^n)$ and $u_t = (u_t^1, \ldots, u_t^m)$ so as to maximize the sum

$$E_s \sum_s^T f(x_t, u_t, \xi_t, t) \quad \text{where } x_t, u_t \in I_t, \text{ subject to the constraints}$$

$$x_{t+1}^j = g^j(x_t, u_t, \xi_t, t), \quad t = s, \ldots, T \quad j = 1, \ldots, n$$

$$\sum_s^T k^j(x_t, u_t, \xi_t, t) = Q^j \quad j = 1, \ldots, J_1$$

$$h^j(x_t, u_t, \xi_t, t) = 0 \quad t = s, \ldots, T \quad j = s, \ldots, J_2$$

$$u_t \in \Gamma_t \in I_t \quad \Gamma_t \text{ some convex subset of } \mathbf{R}^m$$

Of course, $J_1 = 0$ means that there are no summation restrictions, and similarly if $J_2 = 0$. Define the *Hamiltonian*

$$\begin{aligned} H(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) = {} & f(x_t, u_t, \xi_t, t) + \lambda_{t+1} \cdot g(x_t, u_t, \xi_t, t) \\ & + \mu \cdot k(x_t, u_t, \xi_t, t) + \nu_t \cdot h(x_t, u_t, \xi_t, t) \end{aligned}$$

where we of course have used vector notation: $g(\cdot) = (g_1(\cdot), \ldots, g_n(\cdot))$; $\lambda_t = (\lambda_t^1, \ldots, \lambda_t^n)$, etc. The *Lagrangean multipliers* λ^j and ν^j are *stochastic processes* and the $\mu^j : s$ are *stochastic variables* (i.e., independent of t) such that $\mu \in I_s$ and $\lambda_t \in I_t$. It is easy to verify that

$$E_s \sum_s^T f(x_t, u_t, \xi_t, t) = E_s \sum_s^T [H(\cdot) - \lambda_{t+1} \cdot x_{t+1}] - \mu \cdot Q$$

Assumption: *The Hamiltonian $H(x_t, \ldots, t)$ is concave in the variables (x_t, u_t) and continuously differentiable wrt x.*

By lemma 1, if $u_t \in \underset{u \in \Gamma_t}{\operatorname{argmax}} E_t H(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t)$ and $(\tilde{x}_t, \tilde{u}_t)$ is any feasible pair, then

$$\begin{aligned} E_t H(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) & - E_t H(\tilde{x}_t, \tilde{u}_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) \\ & \geq E_t H_x(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) \cdot [x_t - \tilde{x}_t] \end{aligned}$$

Hence, if (x_t, u_t) is a fixed feasible pair such that

$$u_t \in \underset{u \in \Gamma_t}{\operatorname{argmax}} E_t H(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t)$$

and $(\tilde{x}_t, \tilde{u}_t)$ is any feasible pair, then

$$
\begin{aligned}
& E_s \sum_s^T f(x_t, u_t, \xi_t, t) - E_s \sum_s^T f(\tilde{x}_t, \tilde{u}_t, \xi_t, t) \\
& \quad = E_s \sum_s^T \big\{ H(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) - H(\tilde{x}_t, \tilde{u}_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) \big\} \\
& \qquad - E_s \sum_s^T \lambda_{t+1} \cdot [x_{t+1} - \tilde{x}_{t+1}] \\
& \quad = E_s \sum_s^T \big\{ E_t H(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) - E_t H(\tilde{x}_t, \tilde{u}_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) \big\} \\
& \qquad - E_s \sum_s^T \lambda_{t+1} \cdot [x_{t+1} - \tilde{x}_{t+1}] \\
& \quad \geq E_s \sum_s^T E_t H_x(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) \cdot [x_t - \tilde{x}_t] - E_s \sum_s^T \lambda_{t+1} \cdot [x_{t+1} - \tilde{x}_{t+1}] \\
& \quad = E_s \sum_s^T [E_t H_x(x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t, t) - \lambda_t] \cdot [x_t - \tilde{x}_t] \\
& \qquad + \lambda_s \cdot [x_s - \tilde{x}_s] - E_s \big\{ \lambda_{T+1} \cdot [x_{T+1} - \tilde{x}_{T+1}] \big\}
\end{aligned}
$$

The following conditions on $(x_t, u_t, \lambda_t, \mu, \nu_t)$ are obviously sufficient for this expression to be ≥ 0, and hence *sufficient conditions for* (x_t, u_t) *to be a solution to the maximization problem:*

(Ex^j) $\quad E_t H_{x^j}(x_t, u_t, \xi, \lambda_{t+1}, \mu, \nu_t, t) - \lambda_t^j = 0 \quad t = s, \ldots, T \quad j = 1, \ldots, n$

(Mu) $\quad u_t \in \operatorname*{argmax}_{u \in \Gamma_t} E_t H(x_t, u_t, \xi, \lambda_{t+1}, \mu, \nu_t, t) \quad t = s, \ldots, T$

for any $j = 1, \ldots, n$, *either* x_{T+1}^j *is given, or else*
(TUx^j) $\lambda_{T+1}^j = 0$

for any $j = 1, \ldots, n$, *either* x_s^j *is given, or else*
(TLx^j) $\lambda_s^j = 0$

Infinite horizon

Let in the previous case $T = \infty$. The following transversality condition replaces *(TUx^j)* as a sufficient condition, and the derivation parallels that of the case with no uncertainty:

(TU$^\infty$x^j) $\liminf_{t \to \infty} E_\tau \lambda_t \cdot \tilde{x}_t \geq 0$ *for all feasible* $\tilde{x}_t$ *and* $\lim_{t \to \infty} E_\tau \lambda_t \cdot x_t = 0 \ \forall \tau$

The envelope theorem

Let $V = \max E_s \sum_s^T f(x_t, u_t, \xi_t, t)$ and let α be a parameter which does not influence Q, s, T, Γ_t, ξ_t or any of the boundary values of x, if any are given as constraints. Assume also that the probability measures of $x_t, u_t, \xi_t, \lambda_{t+1}, \mu, \nu_t$ conditional on I_t are independent of α, of x_s^j if x_s^j is given as a constraint, and of x_{T+1}^j if x_{T+1}^j is given as a constraint. Assume that the sufficient conditions given above are satisfied and that u_t is uniquely determined by *(Mu)*. Then

$$\frac{dV}{d\alpha} = E_s \sum_s^T H_\alpha(x_t, u_t, \xi, \lambda_{t+1}, \mu, \nu_t, t)$$

$$\frac{dV}{dx_s^j} = \lambda_s^j \quad \textit{if } x_s^j \textit{ is given as a constraint}$$

$$\frac{dV}{dx_{T+1}^j} = -E_s \lambda_{T+1}^j \quad \textit{if } x_{T+1}^j \textit{ is given as a constraint}$$

$$\frac{dV}{dQ_j} = -\mu^j \quad j = 1, \ldots, J_1$$

The proof parallels that of the certainty case, so we omit it.

Bellman's approach to the uncertainty case

A popular way to treat the uncertainty case is to use Bellman's equation. We now show that Bellman's approach often leads to the same equations as Pontryagin's, i.e., those we have derived. Let us look at the maximization problem in the uncertainty case again, where we only have the first type of constraint, i.e., $x_{t+1}^j = g^j(x_t, u_t, \xi_t, t)$, and x_s is given. Let

$$V(x_s; s) \equiv \max E_s \sum_s^T f(x_t, u_t, \xi_t, t)$$

Then *Bellman's principle* states that V satisfies the functional equation

$$V(x_t; t) = \max E_t\{f(x_t, u_t, \xi_t, t) + V(x_{t+1}; t+1)\} \qquad \textit{(FE)}$$
$$\text{where } \; x_{t+1}^j = g^j(x_t, u_t, \xi_t, t)$$

We assume that the probability measure of ξ_t conditional on I_t is independent of u_t and x_t. The first order condition for this maximization problem is then

$$E_t\{f_u(x_t, u_t, \xi_t, t) + V_x(x_{t+1}; t+1) g_u^j(x_t, u_t, \xi_t, t)\} = 0$$

We introduce the notation $\lambda_{t+1} \equiv V_x(x_{t+1}; t+1)$ (cf. the envelope theorem), so this equation becomes

$$E_t\{f_u(x_t, u_t, \xi_t, t) + \lambda_{t+1} g_u^j(x_t, u_t, \xi_t, t)\}$$

which is the same as our equation *(Mu)*: $E_t H_u(x_t, u_t, \xi_t, \lambda_{t+1}, t) = 0$. If we differentiate the functional equation *(FE)* wrt x_t, using the envelope theorem, we get

$$V_x(x_t; t) = E_t\{f_x(x_t, u_t, \xi_t, t) + V_x(x_{t+1}; t+1) g_x^j(x_t, u_t, \xi_t, t)\}$$

which with our λ-notation becomes

$$\lambda_t = E\{f_x(x_t, u_t, \xi_t, t) + \lambda_{t+1} g_x^j(x_t, u_t, \xi_t, t)\}$$

which is precisely our equation *(Ex)*: $E_t H_x(x_t, u_t, \xi_t, \lambda_{t+1}, t) - \lambda_t = 0$.