Exam Part

C8.7 Optimal Control
Mock Exam

May 13, 2025

Do not turn this page until you are told that you may do so

1. In this question we will consider a stochastic optimal control problem, where a controlled process X takes values in $\mathbb{R}$, with controlled dynamics

$$\mathrm{d}X_t = f(t, X_t, U_t)\,\mathrm{d}t + \sigma(t, X_t, U_t)\,\mathrm{d}W_t; \qquad X_0 = x_0.$$

The control U is a progressively measurable process taking values in a topological space $\mathcal{U}$, chosen to minimize the expected cost

$$\mathbb{E}\Big[\int_0^T g(t, X_t, U_t)\,\mathrm{d}t + \Phi(X_T)\Big|\mathcal{F}_t\Big]$$

You may assume whenever needed that all functions are sufficiently integrable and continuous that the state dynamics and expected costs are well defined.

(a) Write down the Hamilton–Jacobi–Bellman equation satisfied by the value function for the control problem, including appropriate boundary conditions.

(b) (i) Suppose you are given a ($C^{1,2}$, polynomial growth) function $v : [0,T] \times \mathbb{R} \to \mathbb{R}$ satisfying $v(\cdot, T) = \Phi(\cdot)$, and a ($C^{1,2}$) function $u^* : [0,T] \times \mathbb{R} \to \mathcal{U}$. If the running cost is written in the form

$$g(t,x,u) = \gamma(t,x,u) - \partial_t v(t,x) - f(t,x,u)\partial_x v(t,x) - \frac{1}{2}\big(\sigma(t,x,u)\big)^2 \partial^2_{xx} v(t,x)$$

for some function γ, give a necessary and sufficient condition on γ such that v is the value function of the control problem, and u^* is the optimal control.

(ii) Write down a control problem where $v(t,x) = \sin(x) + \frac{t}{1+t}$ and $u^*(t,x) = \cos(tx)$, when $\mathcal{U} = \mathbb{R}$, $f(t,x,u) = -u^2 x$, and $\sigma(x) = 2x$.

(c) Now suppose that $f(t,x,u) = g(t,x,u) = 0$, and that Φ is convex. Show that v is convex, and hence describe the optimal strategy.

2. In this question we will consider an undiscounted discrete time, deterministic, finite horizon control problem, where the control takes values in an open set $\mathcal{U}$. We have a continuous state variable X, taking values in $\mathcal{X} = \mathbb{R}$, with initial state $X_0 = x_0$, and one-step controlled dynamics

$$X_{t+1} = X_t + f(t, x, U_t)$$

There is a continuous running cost function $g : \mathbb{T} \times \mathcal{X} \times \mathcal{U}$, and a terminal cost function Φ.

(a) Explain what is meant by a pasting of a control process.

(b) Assuming that f, g and Φ are Lipschitz continuous (uniformly in u), construct an ε-optimal control for this problem, and hence prove the the value function satisfies the Bellman equation.

(c) Using the Bellman equation (which you may assume has a differentiable solution) or otherwise, show that the optimal control satisfies a discrete version of Pontryagin's principle, that is:

$$\begin{aligned} X_{t+1} &= X_t + f(t, X_t, U_t^*); & X_0 &= x_0; \\ q_t &= q_{t+1}\big(1 + \partial_x f(t, X_t, U_t^*)\big) + \partial_x g(t, X_t, U_t^*); & q_T &= \frac{\mathrm{d}\Phi}{\mathrm{d}x}(X_T); \\ 0 &= \partial_u g(t, X_t, U^*) + f(t, X_t, U^*) q_{t+1}. \end{aligned}$$

(d) Explain the key advantages and disadvantages of using Pontryagin's principle, rather than the Bellman equation, to compute the optimal control.

3. Consider an infinite-horizon, discounted, discrete-time, finite-state stochastic control problem. The state has transition dynamics

$$\mathbb{P}(X_{t+1} = x' | X_t = x, U_t = u) = p(x'; x, u)$$

and the goal is to minimize the cost

$$J(U, x) = \sum_{t \geqslant 0} e^{-\rho t} g(X_t, U_t)$$

for a discount rate $\rho > 0$, initial state $X_0 = x$ and cost function g. The controls take values from a compact set, and p and g are both continuous with respect to the control.

(a) [5 marks] Define the Bellman evaluation and optimality operators $\mathcal{T}_u$ and $\mathcal{T}$. Explain how they relate to the cost-to-go $J(U, x)$ and the value function.

(b) [13 marks] (i) Define the standard policy iteration algorithm for solving the control problem.

(ii) Consider the variation of policy iteration where the evaluation procedure remains unchanged, but at each improvement stage a state is chosen uniformly at random, and the action chosen in this state is optimized (and the actions chosen in other states are left unchanged). Show that the value function approximations resulting from this procedure are (componentwise) decreasing, and that they converge to the true value function.
(You should prove any properties of $\mathcal{T}$ and $\mathcal{T}_u$ which you need to use.)

(c) [7 marks] Consider now an agent with a state-dependent discount rate, that is, who wishes to optimize

$$J(U, x) = \sum_{t \geqslant 0} \exp\Big(- \sum_{0 < s \leqslant t} \rho(X_s)\Big) g(X_t, U_t)$$

for some given function $\rho : \mathcal{X} \to (0, \infty)$.

(i) Write down a Bellman equation satisfied by the corresponding value function.

(ii) Show that the Bellman equation admits a bounded solution, which can be approximated by finite-horizon problems.