Sample Exam Questions

Stochastic Decision Processes, January 2004

IHS, Vienna

John Rust

1 . Write down the Bellman equation to solve for the optimal selling strategy of a person trying to sell a house and who must determine an optimal pricing policy to maximize the expected discounted gains from selling it. You may assume that the problem is a finite or infinite horizon decision problem, whichever is easier for you. Assume that the seller has a monetary equivalent reservation value V_l from continuing to live in the house and not selling it, but if the seller can get an offer O that is greater than V_l the seller will prefer to sell rather than continue to live in the house.

A . Assume that the seller is not using a real estate agent but rather places an advertisement in the newspaper. An advertisement can be either for one day only at a cost of $\overline{a}$, or the seller can place an advertisement that runs for an entire week at a daily cost of $\underline{a} < \overline{a}$, reflecting a "quantity discount."

B . Assume that each period at most one buyer can arrive to look at the house and this prospective buyer may or may not decide to make an offer. Assume that the probability that a buyer arrives to view the house is $\lambda(P)$ where P is the seller's advertised price for the house.

C . After viewing the house, the buyer may make an offer. Assume that an offer is made with probability $\pi(P)$ and conditional on making an offer, the actual value of the buyer's offer O is given by a conditional density $f(O|P)$.

D . The seller experiences a monetary equivalent disutility of c for each period the house is being considered for sale due to the need to keep the house clean and to vacate the premises whenever a prospective buyer arrives. If the seller is risk neutral and has discount factor $\beta \in (0, 1)$, write down the Bellman equation for the optimal selling strategy on the part of the seller.

E . From looking at the Bellman equation at this level of generality, how much can you tell about the form of the optimal selling strategy? Is a reservation price strategy generally optimal (i.e. will the sell accept an offer from a buy provided it is above a certain threshold value? If so, will this thresold equal V_l, or will it be higher or lower than this value?)

F. What can you say about the optimal pricing strategy? Will the seller generally start with a high advertised price and decrease it over time, or is it optimal to start with a low price and increase it? How rapidly should the seller change prices: should they be changed every period, or is it optimal to keep the same advertised price for several periods?

G. Finally, generalize the problem to add the decision of whether or not to use a real-estate agent. The cost of using a real-estate agent is that if the house is sold, the seller must pay 6% of the sales proceeds to the real-estate agent. However real estate agents offer a number of benefits to a seller. If the seller uses a real-estate agent, the seller no longer has to pay to advertise the house – the real-estate agent does this as part of the services the agent provides. Most importantly, due to the "connections" the real-estate agent has, the seller can generally get a higher arrival rate of potential buyers. Assume that with a real estate agent, there is a probability $\lambda_r(P)$ that a potential buyer arrives to view the house each period, where $\lambda_r(P) > \lambda(P)$ (i.e. it is higher than when the seller tries to sell the house on his/her own). Furthermore, assume that the real estate agency is able to target its advertisements to a more appropriate set of buyers who have a higher probability of making an offer, $\pi_r(P) > \pi(P)$, and for whom the distribution of offered prices $f_r(O|P)$ stochastically dominates the distribution of offered prices $f(O|P)$ in the absences of a real-estate agent. Write the Bellman equation that determines the optimal selling strategy when a seller can either sell a house on his/her own, or use a real estate agent. Under what conditions might it be optimal to start by selling the house without a real estate agent, and then switch to a real estate agent later, if the seller is unsuccessful in selling the house after a sufficiently long amount of time? Does the possibility of switching from selling a house on your own to using a real estate agent only occur in a finite horizon formulation of the selling problem, or could it happen in the infinite horizon problem as well?

2. Derive the McQueen-Porteus Error bounds on successive iterations of the Bellman equation.

3. Prove that the fixed point to a contraction mapping on a Banach space exists and is unique.

4 . Consider an infinite horizon dynamic programming problem with utility function given by

$$u(s,a) = \lambda_2 a^2 + \lambda_1 a + \lambda_0 + s[\rho_0 + \rho_1 a] + \mu s^2$$

and transition probability

$$p(s'|s,a) = \frac{1}{\sqrt{2\pi}\sigma}\, exp\left\{-(s' - \kappa_0 - \kappa_1 a - bs)^2/(2\sigma^2)\right\}$$

Assume the decision maker wants to maximize expected discounted utility over an infinite horizon with discount factor $\beta \in (0,1)$.

A . Write the Bellman equation for this problem.

B . Show that the value function $V(s)$ is a quadratic function of s and the optimal decision rule $\alpha(s)$ is a linear function of s.

C . Does the Bellman equation have only 1 solution in this (unbounded utility) case, or is it possible that it has multiple solutions?

5 . Consider the problem of integrating a function $f(x)$ on the $[0,1]$ interval. We want to characterize the optimal *deterministic* algorithm for integration.

A . Suppose that our algorithm consists of evaluating the function at a finite number of points $(x_1,\ldots,x_n)$ in $[0,1]$ and then using the resulting values $(f(x_1),\ldots,f(x_n))$ to form an estimate of the integral. Provide an example of an integration rule that works this way. How should the *information* $I = (f(x_1),\ldots,f(x_n))$ on the function f be used?

B . Suppose that all we knew about the function f (besides being able to evaluate it any any finite number of points in $[0,1]$) is that it is continuous. Show that in the worst case, the error in trying to integrate the function f using only its values at a finite number of points $(f(x_1),\ldots,f(x_n))$ is ∞, i.e.

$$\sup_{f\in\mathcal{F}} \left| \int_0^1 f(x)dx - \phi(f(x_1),\ldots,f(x_n)) \right| = \infty$$

where ϕ is some *algorithm* for using the *information on* f $I = (f(x_1),\ldots,f(x_n))$ to compute an approximation to the integral of f.

C . Thus, we conclude that from a worst case error basis, it is impossible to integrate functions if all that we know about them is that they are continuous. More prior information on the smoothness and boundedness of the class of functions $\mathcal{F}$ is required. Now consider what the worst case error would be if the class we are integrating over is the class $\mathcal{F}_L$ of *Lipschitz continuous functions on* $[0,1]$ *with uniform Lipschitz bound* L. This is the class of functions $\mathcal{F}_L$ satisfying

$$|f(x) - f(x)| \leq L|x-y|, \quad x, y \in [0,1], \quad f \in \mathcal{F}_L.$$

D . An *algorithm* for computing the integral of a function $f \in \mathcal{F}$ can be written as a composition of two functions, $\phi_n : R^n \rightarrow R$ and $I_n : [0,1]^n \rightarrow R^n$ where $I_n = (f(s_1), \ldots, f(s_n))$ is the *information* (or sample) on the unknown function f we wish to integrate and ϕ_n is an algorithm that combines this information into an estimate of the integral:

$$\int_0^1 f(x)dx \approx \phi_n(f(s_1), \ldots, f(s_n)).$$

Thus, we are seeking a rule for choosing the sample points $(s_1, \ldots, s_n)$ and the function ϕ_n that minimize the worst case integration error:

$$r(n) = \underset{s_1,\ldots,s_n}{inf} \; \underset{\phi_n}{inf} \; \underset{g \in \mathcal{F}(f(s_1),\ldots,f(s_n))}{sup} \left| \phi_n(f(s_1), \ldots, f(s_n)) - \int_0^1 g(x)dx \right|,$$

where

$$\mathcal{F}(f(s_1), \ldots, f(s_n)) = \left\{ g \in \mathcal{F} | g(s_1) = f(s_1), \ldots, g(s_n) = f(s_n) \right\}.$$

Thus, $\mathcal{F}(f(s_1), \ldots, f(s_n))$ is the equivalence class of functions in $\mathcal{F}$ that have the same information (i.e. have the same values over the n sample points $(s_1, \ldots, s_n)$) as the true function f that we are trying to integrate. Since we assume that we don't know the true f at all points but only at the n points $(s_1, \ldots, s_n)$, we consider the worst case error by computing the function $g \in \mathcal{F}(f(s_1), \ldots, f(s_n))$ whose actual integral is as far away as possible from the approximate integral $\phi_n(f(s_1), \ldots, f(s_n))$.

E . Show that $\mathcal{F}(f(s_1), \ldots, f(s_n))$ is a set of functions in $\mathcal{F}$ bounded above by an *upper envelope* $\overline{f}$ and and a *lower envelope* $\underline{f}$, and that $\overline{f}$ and $\underline{f}$ are piecewise-linear functions with slopes everywhere equal to $\pm L$ that satisfy

$$\underline{f}(s_i) = \overline{f}(s_i) = f(s_i), \quad i = 1, \ldots, n.$$

F. Show that the optimal algorithm ϕ_n^* is given by

$$\phi_n^*(f(s_1), \ldots, f(s_n)) = \int_0^1 f_{\text{mid}}(x)dx,$$

where $f_{\text{mid}} = (\underline{f} + \overline{f})/2$.

G. Let f_{pwl} be the piecewise linear interpolant of the points $(0, f(s_1))$, $(s_1, f(s_1))$, $\ldots, (s_n, f(s_n)), (1, f(s_n))$. Show that

$$\int_0^1 f_{\text{mid}}(x)dx = \int_0^1 f_{\text{pwl}}(x)dx.$$

H. Show that

$$\int_0^1 f_{\text{pwl}}(x)dx = f(s_1)s_1 + \sum_{i=1}^{n-1} \frac{1}{2}(f(s_i) + f(s_{i+1})(s_{i+1} - s_i) + f(s_n)(1 - s_n).$$

Thus, the optimal integration algorithm for the class $\mathcal{F}$ is the *modified trapezoidal rule.*

I. Show that there is no loss of generality in restricting attention to the special case of *zero information,* i.e. where $f(s_1) = f(s_2) = \cdots = f(s_n) = 0$, i.e. show that

$$\sup_{g \in \mathcal{F}(f(s_1),\ldots,f(s_n))} \left| \phi_n^*(f(s_1), \ldots, f(s_n)) - \int_0^1 g(x)dx \right| =$$

$$\sup_{g \in \mathcal{F}(0,\ldots,0)} \left| \phi_n^*(0, \ldots, 0) - \int_0^1 g(x)dx \right| =$$

$$\sup_{g \in \mathcal{F}(0,\ldots,0)} \left| \int_0^1 g(x)dx \right|.$$

J. Show that

$$\sup_{g \in \mathcal{F}(0,\ldots,0)} \left| \int_0^1 g(x)dx \right| = L \left(\frac{1}{2}s_1^2 + \frac{1}{4}\sum_{i=1}^{n-1} (s_{i+1} - s_i)^2 + \frac{1}{2}(1 - s_n)^2 \right).$$

K. Using calculus, derive the optimal placement of the sample points. Show that the optimal points satisfy

$$s_i^* = \frac{2i - 1}{2n}$$

show that the worst case error bound using the optimal integration algorithm and the optimally placed points satisfies

$$r(n) = \frac{L}{4n}.$$

L. Show that for the optimally chosen points, the optimal integration algorithm takes the form of a *quasi monte carlo algorithm:*

$$\phi_n^*(f(s_1^*), \ldots, f(s_n^*)) = \frac{1}{n} \sum_{i=1}^{n} f(s_i^*).$$

How does this result compare with the quasi monte carlo algorithm for the minimum discrepancy points $(t_1^*, \ldots, t_n^*)$ that you derive in problem 6 below?

6. The *discrepancy* of a set of n points $(t_1, \ldots, t_n)$ in $[0,1]^d$ (the d-dimensional hypercube) is given by

$$D_n^*(t_1, \ldots, t_n) = \sup_{B \in \mathcal{B}} |\lambda_n(B) - \lambda(B)|,$$

where $\mathcal{B}$ is the set of all *normalized subrectangles* of $[0,1]^d$,

$$\mathcal{B} = \left\{ B \subset [0,1]^d | B = \prod_{i=1}^{d} [0, b_i], \quad b_i \in [0,1] \right\},$$

and $\lambda(B)$ is the *Lebesgue measure* of B,

$$\lambda(B) = \prod_{i=1}^{d} b_i,$$

and $\lambda_n(B)$ is the *empirical measure* of B,

$$\lambda_n(B) = \frac{1}{n} \sum_{i=1}^{n} I\{s_i \in B\}.$$

where $I\{s \in B\}$ is the indicator function,

$$I\{s \in B\} = \begin{cases} 1 & \text{if} \quad s \in B \\ 0 & \text{otherwise.} \end{cases}$$

A. Consider the one dimensional case, $d = 1$. Find a formula for the discrepancy, $D_n^*(t_1, \ldots, t_n)$.

B. Find a formula for the *minimal discrepancy* points $(t_1^*, \ldots, t_n^*)$, i.e. the points that solve

$$(t_1^*, \ldots, t_n^*) = \underset{(t_1, \ldots, t_n)}{argmin} \, D_n^*(t_1, \ldots, t_n).$$