Due Date: This assignment is due on Thursday, 4 June, 2009(with the natural 1 day extension for SCPD students), by 5pm in the box outside Durand 112. The LaTeX incentive policy is in effect.

Problem 1: In this problem, you will provide an alternative proof of the fact that finite state Markov chains typically converge to equilibrium exponentially rapidly. In particular, consider an irreducible finite state Markov chain having a transition matrix of the form

$$P = \alpha\Lambda + (1-\alpha)\,Q$$

for some $\alpha \in (0,1]$. where Λ is a rank one stochastic matrix with identical rows.

1. Argue that if P has a strictly positive column, it is of the above form.
2. Prove that $Q\Lambda = \Lambda^2 = \Lambda$.
3. Using the equilibrium calculated in Problem 7 of Assignment 4, compute

$$P^n - \Pi$$

in terms of α, Λ, and Q, where Π is the rank one matrix having all rows equal to the equilibrium distribution π.

4. Use matrix norms to show that there exists $a > 0$ such that $\exp(an)\,\|P^n - \Pi\| \to 0$ as $n \to \infty$.

Solution:

1. Assume $\Lambda = \begin{bmatrix} \lambda^T \\ \vdots \\ \lambda^T \end{bmatrix}$ and $P = \begin{bmatrix} p_1^T \\ \vdots \\ p_n^T \end{bmatrix}$. If P has identical rows, then $P = \Lambda$ with $\alpha = 1$. Otherwise $0 \le \alpha < 1$, and we have the following expression:

$$P = \alpha \begin{bmatrix} \lambda^T \\ \vdots \\ \lambda^T \end{bmatrix} + (1-\alpha) \begin{bmatrix} \frac{1}{1-\alpha}p_1^T - \frac{\alpha}{1-\alpha}\lambda^T \\ \vdots \\ \frac{1}{1-\alpha}p_n^T - \frac{\alpha}{1-\alpha}\lambda^T \end{bmatrix}$$

So the rest is to find the suitable α and λ. We require that (i)$\lambda \ge 0$ (ii)the sum of all elements of λ is 1 (iii) $\frac{1}{1-\alpha}p_i - \frac{\alpha}{1-\alpha}\lambda \ge 0$, and (iv) the sum of all elements of $\left(\frac{1}{1-\alpha}p_i - \frac{\alpha}{1-\alpha}\lambda\right)$ is 1. Note that (iv) is automatically true if (ii) is true. Furthermore, (iii) is equivalent to that $p_i \ge \alpha\lambda$. From this condition a good guess of λ is e_k, the vector with k-th entry 1 and 0 otherwise. Now let k be the location of the strictly positive column in P and α be the minimum value in this column. Then conditions (i)-(iv) are satisfied by this setting.

2. Just note that for any row vector v^T whose the sum of all elements is 1, $v^T\Lambda = \Lambda$. Then $Q\Lambda = \Lambda^2 = \Lambda$.

3. By the result from HW4

$$\begin{aligned} P^n - \Pi &= [\alpha\Lambda + (1-\alpha)Q]^n - \begin{bmatrix} \alpha\lambda^T\left(I - (1-\alpha)\,Q\right)^{-1} \\ \vdots \\ \alpha\lambda^T\left(I - (1-\alpha)\,Q\right)^{-1} \end{bmatrix} \\ &= [\alpha\Lambda + (1-\alpha)Q]^n - \alpha\Lambda\left(I - (1-\alpha)\,Q\right)^{-1} \\ &= [\alpha\Lambda + (1-\alpha)Q]^n - \alpha\Lambda\sum_{i=0}^{\infty}(1-\alpha)^i\,Q^i \end{aligned}$$

Use part 2 and induction to argue that

$$[\alpha\Lambda + (1-\alpha)Q]^n = \alpha\Lambda \sum_{i=0}^{n-1}(1-\alpha)^i Q^i + (1-\alpha)^n Q^n$$

This formula holds for $n = 1$ and suppose it also holds for some n. Then

$$\begin{aligned}
[\alpha\Lambda + (1-\alpha)Q]^{n+1} &= [\alpha\Lambda + (1-\alpha)Q]\left[\alpha\Lambda \textstyle\sum_{i=0}^{n-1}(1-\alpha)^i Q^i + (1-\alpha)^n Q^n\right] \\
&= [\alpha\Lambda + (1-\alpha)Q]\,\alpha\Lambda \textstyle\sum_{i=0}^{n-1}(1-\alpha)^i Q^i \\
&\quad + \alpha(1-\alpha)^n \Lambda Q^n + (1-\alpha)^{n+1} Q^{n+1} \\
&= \alpha\Lambda \textstyle\sum_{i=0}^{n-1}(1-\alpha)^i Q^i + \alpha(1-\alpha)^n \Lambda Q^n + (1-\alpha)^{n+1} Q^{n+1} \\
&= \alpha\Lambda \textstyle\sum_{i=0}^{n}(1-\alpha)^i Q^i + (1-\alpha)^{n+1} Q^{n+1}
\end{aligned}$$

Finally,

$$P^n - \Pi = (1-\alpha)^{n+1} Q^{n+1} - \alpha\Lambda \sum_{i=n}^{\infty} (1-\alpha)^i Q^i.$$

4. From the computation of part 3,

$$\begin{aligned}
\|P^n - \Pi\| &= \|(1-\alpha)^{n+1} Q^{n+1} - \alpha\Lambda \textstyle\sum_{i=n}^{\infty} (1-\alpha)^i Q^i\| \\
&\leq \|(1-\alpha)^{n+1} Q^{n+1}\| + \|\alpha\Lambda \textstyle\sum_{i=n}^{\infty} (1-\alpha)^i Q^i\| \\
&\leq (1-\alpha)^{n+1}\|Q\|^{n+1} + \alpha\|\Lambda\| \textstyle\sum_{i=n}^{\infty} (1-\alpha)^i \|Q\|^i
\end{aligned}$$

Take the norm to be the infinity norm. Then $\|\Lambda\|_\infty$ and $\|Q\|_\infty$ are equal to 1. That implies

$$\begin{aligned}
\|P^n - \Pi\|_\infty &\leq (1-\alpha)^{n+1} + \alpha \textstyle\sum_{i=n}^{\infty} (1-\alpha)^i \\
&\leq (1-\alpha)^{n+1} + \alpha \frac{(1-\alpha)^n}{1-(1-\alpha)} \\
&\leq 2(1-\alpha)^n \\
&= 2e^{n \log(1-\alpha)}
\end{aligned}$$

Let $0 < a < -\log(1-\alpha)$, $e^{an}\|P^n - \Pi\|_\infty \to 0$ as $n \to \infty$. Since all norms are equivalent in the finite dimensional space, this result holds for any matrix norms.

Problem 2: In this problem, we will use Lyapunov method, to study equilibrium behavior for state space models. Let $(X_n : n \geq 0)$ be an $\mathbb{R}^d$-valued stochastic sequence defined by

$$X_{n+1} = FX_n + Z_{n+1},$$

where $(Z_i : i \geq 1)$ is an iid sequence of $\mathbb{R}^d$-valued random vectors having a continuous and positive density and F is a matrix having spectral radius less than one.

1. Compute

$$P_x(X_1 \leq y)$$

in terms of the distribution of Z_1.

2. Compute the conditional density of X_1 (i.e. the transition density), namely the function $p(x, y)$ such that

$$P_x(X_1 \in B) = \int_B p(x, y)\, dy.$$

3. Prove that for any compact set K,

$$\inf_{x\in K} p(x, y) > 0.$$

4. Prove that if $\mathrm{E}\|Z_1\| < \infty$, then $g(x) = \|x\|$ is a Lyapunov function for $(X_n : n \geq 0)$. (Hence, according to Theorem 8.15 of the Course Notes, X has an equilibrium distribution π.)
   Note: $\|\cdot\|$ is not a 2-norm in general, and please specify the norm you use.

5. (Extra Credit). Prove that if $\mathrm{E}\log(1 + \|Z_1\|) < \infty$, Theorem 8.15 can still be applied. (You must modify the Lyapunov function.)

Solution:

1. Assume f is the density function of Z_i's.

$$\begin{aligned} P_x(X_1 \leq y) &= P_x(FX_0 + Z_1 \leq y) = P(Fx + Z_1 \leq y) \\ &= P(Z_1 \leq y - Fx) = \int_{z \leq y - Fx} f(z)\,dz \\ &= \int_{z \leq y} f(z - Fx)\,dz \end{aligned}$$

2. By the part 1,

$$\begin{aligned} P_x(X_1 \in B) &= P(Fx + Z_1 \in B) = P(Z_1 \in B - Fx) \\ &= \int_{B - Fx} f(z)\,dz = \int_B f(z - Fx)\,dz \end{aligned}$$

   Therefore $p(x, y) = f(y - Fx)$.

3. Note that f is continuous and strictly positive in $\mathbb{R}^d$. Then for a fixed y, $p(x, y) = f(y - Fx)$ is also continuous and strictly positive for all $x \in \mathbb{R}^d$. So $p(x, y)$ can reach its minimum in any compact K, i.e. $\inf_{x \in K} p(x, y) = p(x^*, y) > 0$, where $x^* \in K$ is the minimizer.

4. By theorem 8.15, we have to prove

   (a) $E_x\|X_1\| \leq \|X_1\| - \epsilon$ for all $x \in A^c$

   (b) $\sup_{x \in A} E_x\|X_1\| < \infty$

   (c) There is m such that $P_x(X_m \in dy) \geq \lambda\phi(y)$, for all $x \in A$

   To prove (a),

$$\begin{aligned} E_x\|X_1\| &= \int_{\mathbb{R}^d} \|Fx + y\| f(y)\,dy \\ &\leq \int_{\mathbb{R}^d} \|Fx\| f(y)\,dy + \int_{\mathbb{R}^d} \|y\| f(y)\,dy \\ &= \|Fx\| + E\|Z_1\| \\ &\leq \|F\|\|x\| + E\|Z_1\| \end{aligned}$$

   Since $\rho(A) < 1$, we can find a specific norm $\|\cdot\|$ (depending on F) such that $\|F\| = \alpha < 1$. Then $\exists\epsilon > 0$ such that $E_x\|X_1\| < \|x\| - \epsilon$ for $\|x\|$ sufficient large.

   To prove (b), note that from (a) we have $E_x\|X_1\| \leq \alpha\|x\| + E\|Z_1\|$. So $\sup_{x \in A} E_x\|X_1\| \leq \sup_{x \in A} \alpha\|x\| + E\|Z_1\| < \infty$ if we choose A to be a bounded set.

   To prove (c), from part 2 and part 3, we have $P_x(X_1 \in dy) = p(x, y)$ and $\inf_{x \in K} p(x, y) > 0$. We want $p(x, y) \geq \lambda\phi(y)$, i.e. $\lambda\phi(y) = \inf_{x \in A} p(x, y)$. So $\phi(y) = \frac{\inf_{x \in A} p(x,y)}{\int \inf_{x \in A} p(x,y)dy}$ and $\lambda = \int \inf_{x \in A} p(x, y)\,dy$. If A is a compact set, then $\inf_{x \in A} p(x, y) > 0$ for all $y \Rightarrow \lambda > 0$ and $\phi(y)$ is well-defined.

   By above argument, we need A to be a compact set and to satisfy the condition (a). Finally $g(x) = \|x\|$, which is the norm in (a), is a stochastic Lyapunov function.

5. It suffices to find another stochastic Lyapunov function $h(x)$. Now we claim that $h(x) = \log(2 + \|x\|)$. Since (c) is unchanged, we only have to check (a) (b). Here we use the fact that for nonnegative a and b, $(1 + a) + (1 + b) \leq 2(1 + a)(1 + b)$.

$$\begin{aligned} E_x \log(2 + \|X_1\|) &\leq E_x \log(2 + \|Z_1\| + \alpha\|x\|) \\ &\leq E\log(1 + \|Z_1\|) + \log 2 + \log(1 + \alpha\|x\|) \\ &\leq \log(2 + \|x\|) - \epsilon \\ &= h(x) - \epsilon \end{aligned}$$

for $\|x\|$ sufficient large. Also, $\sup_{x\in A} E_x \log(2+\|X_1\|) \le E\log(1+\|Z_1\|)+\log 2+\sup_{x\in A}\log(1+\alpha\|x\|) < \infty$ if A is bounded. Like 4., choose A to be a compact set and to satisfy (a). We conclude that $h(x)$ is a stochastic Lyapunov function.

Problem 3: (Heyman and Sobel, Exercise 3-16) Consider a small commercial catfish farming pond. Harvests occur annually and the harvest takes only a few days. The farmer can accurately estimate the amount (tons) of fish in the pond and then decide on the harvest quantity. The farmer's experience and U.S. Department of Agriculture data lead to the following probabilities:

Tons of fish remaining after harvest this year	Tons of fish at harvest time next year: 0	1	2	3	4	5
0	1	0	0	0	0	0
1	0	0.9	0.1	0	0	0
2	0	0	0.8	0.1	0.1	0
3	0	0	0	0.7	0.2	0.1
4	0	0	0	0	0.6	0.4
5	0	0	0	0	0	1

The farmer's annual discount factor is 0.9 and the net profit (i.e., revenue - cost) is \$500 per ton. The farmer is young and optimistic and use an infinitely long planning horizon. There are 5 tons in the pond at the present harvest time. How many tons should be harvested now to maximize the expected present value of the net profit? What will the expected present value of the profit be when an optimal policy is used? Justify your answer.

Solution: Let a be the number of tons harvested. We set this up as a linear program:

$$\begin{aligned} \text{minimize} \quad & \sum_x v(x) \\ \text{s.t.} \quad & v(x) \ge 500a + .9\sum_y p_a(x,y)v(y), \quad x\in\mathbb{S}, a\in\mathcal{A}(x) \end{aligned}$$

This can then be solved with `linprog` in MATLAB. In particular, we are not so much concerned with the solution which would give us the value function but rather which constraints are active. In MATLAB, we write

```
>> [X,FVAL,EXITFLAG,OUTPUT,LAMBDA] = linprog(ones(1,6),A,b)
```

where `A` is the inequality constraint matrix and `b` is the inequality constraint vector. We then use

```
LAMBDA.ineqlin
```

to see which Lagrange multipliers are non-zero and thus which constraints are active. The optimal actions for each state are $a_0 = 0$, $a_1 = 0$, $a_2 = 0$, $a_3 = 1$, $a_4 = 2$, and $a_5 = 3$. The optimal value function is $v(0) = 0$, $v(1) = \$639.50$, $v(2) = \$1,350$, $v(3) = \$1,850$, $v(4) = \$2,350$, and $v(5) = \$2,850$.

Problem 4: Consider a system that represents a conventional ("circuit-switched") network. The network has R routes. Each route connects a source node to a destination node along a set of links. The i-th link has the capacities to simultaneously handle c_i calls. Customers attempt to initiate calls on route r at rate λ_r $(1 \le r \le R)$. A call is accepted on route r only if there is sufficient link capacity on each link along the route. If insufficient capacity exists somewhere along the route, the customer's call is dropped (i.e. the customer gets a busy signal). A customer that successfully places a call completes the call at rate γ_r, at which time all the "circuit links" are released.

We model this system as a network through an "incidence matrix" $A = (A_{ij} : 1 \le i \le m, 1 \le r \le R)$, with $A_{ir} = 1$ if route r uses link i and is 0 otherwise (where m is the number of links). If we model the network as a Markov jump process $X(t) = (X_1(t), \ldots, X_R(t))^T$, the state space is $\mathbb{S} = \{n = (n_1, \ldots, n_R)^T \in \mathbb{Z}_+^R : A \cdot n \le c\}$, where $c = (c_1, \ldots, c_m)^T$.

1. Specify the rate matrix Q for this model.
2. Show that the equilibrium distribution $\pi = (\pi(n) : n \in \mathbb{S})$ is given by
$$\pi(n) = \frac{\nu(n)}{\sum_{m \in \mathbb{S}} \nu(m)},$$
where
$$\nu(m) = \prod_{r=1}^{R} \frac{(\lambda_r/\gamma_r)^{m_r}}{m_r!}.$$
3. Now, suppose that we look at the asymptotic setting in which the arrival rates are large and capacities are large (which is typical of the real world). We let $\lambda_r = \hat{\lambda}_r s$ and $c_i = \hat{c}_i s$ and let $s \to \infty$. Let n^* be the mode of $(\pi(n) : n \in \mathbb{S})$, so that $\pi(n^*) \ge \pi(n)$ for $n \in \mathbb{S}$ (i.e. n^* is the "most likely point"). Argue heuristically that
$$n* \approx sx^*,$$
where x^* is the solution of a convex optimization problem. (Hint: you will need to use Sterling's approximation.)

   Remark: This scaling limit is sometimes called the "thermodynamic limit". It can be shown that if $X(\infty)$ is the equilibrium of $(X(t) : t \ge 0)$, then $s^{-1}X(\infty) \Rightarrow x^*$ as $s \to \infty$.

Solution:

1. Note that $X(t) = (X_1(t), \ldots, X_R(t))^T$, and $X_i(t)$ is the number of customers on the route i at time t. The rate matrix $Q = (q_{i,j})$, where $i, j \in \mathbb{S}$. Therefore,

   $q_{n,n+e_k} = \lambda_k$, if $A(n+e_k) \le c$.

   $q_{n,n-e_k} = n_k\gamma_k$, if $n - e_k \ge 0$.

   $q_{n,n} = -\left(\sum_{k:A(n+e_k)\le c} \lambda_k + \sum_{k:n-e_k\ge 0} n_k\gamma_k\right)$.

   $q_{i,j} = 0$, otherwise.

   Also we could derive

   $q_{n+e_k,n} = (n_k+1)\gamma_k$ if $A(n+e_k) \le c$.

   $q_{n-e_k,n} = \lambda_k$ if $n - e_k \ge 0$.

2. Obviously, $\sum_{n\in\mathbb{S}} \pi(n) = 1$ and $\pi \ge 0$. Thus we only need to check $\sum_{k\ne n} \pi(k) q_{k,n} = -\pi(n) q_{n,n}$ for all $n \in \mathbb{S}$, and equivalently, $\sum_{k\ne n} \nu(k) q_{k,n} = -\nu(n) q_{n,n}$. That means
$$\sum_{k:A(n+e_k)\le c} \nu(n+e_k) q_{n+e_k,n} + \sum_{k:n-e_k\ge 0} \nu(n-e_k) q_{n-e_k,n} = -\nu(n) q_{n,n}$$
   Use $\nu(m) = \Pi_r \frac{(\lambda_r/\gamma_r)^{m_r}}{m_r!}$ and part 1.
$$\sum_{k:A(n+e_k)\le c} \prod_r \frac{(\lambda_r/\gamma_r)^{(n+e_k)_r}}{(n+e_k)_r!} (n_k+1)\gamma_k + \sum_{k:n-e_k\ge 0} \prod_r \frac{(\lambda_r/\gamma_r)^{(n-e_k)_r}}{(n-e_k)_r!} \lambda_k$$
$$= \prod_r \frac{(\lambda_r/\gamma_r)^{n_r}}{n_r!} \left(\sum_{k:A(n+e_k)\le c} \lambda_k + \sum_{k:n-e_k\ge 0} n_k\gamma_k\right)$$

Fix e_k such that for $A(n+e_k) \le c$,

$$\prod_r \frac{(\lambda_r/\gamma_r)^{(n+e_k)_r}}{(n+e_k)_r!}(n_k+1)\gamma_k = \prod_r \frac{(\lambda_r/\gamma_r)^{n_r}}{n_r!}\lambda_k$$

is true. Fix e_k such that for $n - e_k \ge c$,

$$\prod_r \frac{(\lambda_r/\gamma_r)^{(n-e_k)_r}}{(n-e_k)_r!}\lambda_k = \prod_r \frac{(\lambda_r/\gamma_r)^{n_r}}{n_r!}n_k\gamma_k$$

is also true. Then the proof is completed.

3. Now we're solving a constrained optimization problem:

$$\max_{n\ge 0,\, An\le c} \pi(n)$$

and equivalently,

$$\max_{n\ge 0,\, An\le c} \nu(n) = \max_{n\ge 0,\, An\le c} \Pi_r \frac{(\lambda_r/\gamma_r)^{n_r}}{n_r!}$$

Use log function to simplify it:

$$\begin{aligned}\max_{n\ge 0,\, An\le c} \log\nu(n) &= \max_{n\ge 0,\, An\le c} \sum\nolimits_r \left(n_r \log\tfrac{\lambda_r}{\gamma_r} - \log n_r!\right)\\ &= \max_{n\ge 0,\, An\le s\hat{c}} \sum\nolimits_r \left(n_r \log\tfrac{s\hat{\lambda}_r}{\gamma_r} - \log n_r!\right)\end{aligned}$$

Heuristically, if $s \to \infty$ then $n_r \to \infty$ as well. So we can use the Sterling approximation: $n! \approx \sqrt{2\pi n}\left(\frac{n}{e}\right)^n$. Instead we solve:

$$\begin{aligned}&\max_{n\ge 0,\, An\le s\hat{c}} \sum\nolimits_r \left(n_r \log\tfrac{s\hat{\lambda}_r}{\gamma_r} - \log\sqrt{2\pi n_r}\left(\tfrac{n_r}{e}\right)^{n_r}\right)\\ =&\max_{n\ge 0,\, An\le s\hat{c}} \sum\nolimits_r \left(n_r \log\tfrac{s\hat{\lambda}_r}{\gamma_r} - \log\sqrt{2\pi} - \tfrac{1}{2}\log n_r - n_r\log n_r + n_r\right)\end{aligned}$$

After simplification,

$$\begin{aligned}&\max_{n\ge 0,\, An\le s\hat{c}} \sum\nolimits_r \left(n_r \log\tfrac{s\hat{\lambda}_r}{\gamma_r n_r} - \tfrac{1}{2}\log n_r + n_r\right)\\ =&\max_{\frac{n}{s}\ge 0,\, A\frac{n}{s}\le \hat{c}} \sum\nolimits_r s\left(\tfrac{n_r}{s}\log\tfrac{\hat{\lambda}_r}{\gamma_r\frac{n_r}{s}} - \tfrac{1}{2s}\log n_r + \tfrac{n_r}{s}\right)\end{aligned}$$

Let $x = \frac{n}{s}$, then the minimizer $n^* = sx^*$, where x^* is the maximizer of

$$\max_{x\ge 0,\, Ax\le \hat{c}} \sum_r \left(x_r \log\frac{\hat{\lambda}_r}{\gamma_r x_r} - \frac{1}{2s}\log s x_r + x_r\right)$$

Note that when s is sufficient large, $-\frac{1}{2s}\log s x_r$ is almost zero because x_r is bounded, so this problem can be well-approximated by

$$\max_{x\ge 0,\, Ax\le \hat{c}} \sum_r \left(x_r \log\frac{\hat{\lambda}_r}{\gamma_r x_r} + x_r\right)$$

Let $f_r(x_r) = x_r \log\frac{\hat{\lambda}_r}{\gamma_r x_r} + x_r$. Then $f_r''(x_r) = -\frac{1}{x_r} < 0$ for $x_r > 0$, i.e. x^* is the maximizer of the concave maximization problem or equivalently, the solution of the standard convex minimization problem:

$$\min_{x\ge 0,\, Ax\le \hat{c}} -\sum_r \left(x_r \log\frac{\hat{\lambda}_r}{\gamma_r x_r} + x_r\right)$$

Problem 5: Suppose that $X = (X(t) : t \geq 0)$ is a Markov jump process representing the price of an asset. We have an option on the asset under which we collect $\exp(-\alpha T) r(X(T))$ if we exercise the option at time T.

Provide the HJB equation which characterizes the value function for computing

$$\sup_T \mathbf{E}_x[\exp(-\alpha T) r(X(T))]$$

and describe the optimal exercise time T^* in terms of the value function.

Solution: Let

$$V(x) = \sup_T \mathbf{E}_x[\exp(-\alpha T) r(X(T))].$$

If we set T^* as the optimal stopping rule, then for any arbitrary T, we have

$$V(x) \geq \mathbf{E}_x[\exp(-\alpha T) r(X(T))].$$

In particular, for $T = 0$ we have $V(x) \geq r(x)$ and for $T = \infty$, $V(x) \geq 0$ (this corresponds to never stopping and therefore reaping no reward). Now let's consider the time $T = h + T^*$ for a small time h. This corresponds to first waiting h time units, and then using T^* after time h. By "small time h analysis",

$$\begin{aligned}
V(x) &\geq \mathbf{E}_x[\exp(-\alpha(h + T^*)) r(X(h + T^*))] \\
&= \mathbf{E}_x[\exp(-\alpha h) \mathbf{E}_{X(h)}[\exp(-\alpha(h + T^*)) r(X(T^*))]] && \text{At } h\text{, we execute } T^* \\
&= \mathbf{E}_x[\exp(-\alpha h) V(X(h))] \\
&= (1 - \alpha h) \mathbf{E}_x[V(X(h))] + o(h).
\end{aligned}$$

We turn to computing $\mathbf{E}_x[V(X(h))]$. By the definition of the rate matrix Q,

$$p(h, x, y) = q_{x,y} h + o(h), \quad p(h, x, x) = 1 - q_{x,x} h + o(h)$$

where $q_{x,y}$ are entries in Q. Thus,

$$\mathbf{E}_x[V(X(h))] = ((I + Qh)V)(x) + o(h).$$

Plugging this in above, we have

$$V(x) \geq (1 - \alpha h)(V(x) + (QV)(x)h + o(h)),$$

which, when we rearrange, divide by h, and take $h \downarrow 0$, we have

$$QV - \alpha V \geq 0.$$

The HJB equation for this problem is

$$\begin{aligned}
QV + \alpha V &\geq 0 \\
V &\geq r \\
V &\geq 0, \quad \forall x \in \mathbb{S}
\end{aligned}$$

This is also known as a variational inequality. We are looking for the minimal function that satisfies these inequalities. Note the similarity to turning the HJB equation into a linear program as was the case in the infinite horizon control.

The question is, how do we characterize the optimal stopping rule? We get this from the similar rule for the discrete time case. Namely, we are looking for the time when the value function V is the same as the reward function. If you consider the stopping rule as a game, you see that as time evolves, you always have the option to stop, but you won't stop if the future holds a higher payout. Thus, the optimal stopping time is characterized by $T^* = \inf\{t \geq 0 : r(X(t)) = V(X(t))\}$. In the context of the inequalities above, we see that for any x such that $V(x) = r(x)$, this constitutes the "stopping region." For any x such that $V(x) > r(x)$, this constitutes the continuation region.

To conclude, the value function is the minimal solution to

$$\begin{aligned} QV + \alpha V &\geq 0 \\ V &\geq r \\ V &\geq 0, \quad \forall x \in \mathbb{S} \end{aligned}$$

and the optimal stopping rule is defined as $T^* = \{t \geq 0 : X(t) \in C^c\}$ where $C = \{x : V(x) > r(x)\}$ and $C^c = \{x : r(x) = V(x)\}$.