Lawler (1995)
Chapter 4
Optimal Stopping

4.1 Optimal Stopping of Markov Chains

- Suppose P is the transition matrix for a discrete-time Markov chain X_n with a finite state space S
  - Let $f(x)$ denote the **payoff function** telling the payoff if the chain is stopped in state x
  - A **stopping time** (or **stopping rule**) T is a random variable that gives the time at which the chain is stopped
    - Stopping time T takes values in the set $\{0,1,\ldots\}$
    - Stopping time T should be such that the decision at time n must be based only on what has happened up through step n. In other words: $I\{T = n\}$ is measurable with respect to $X_1,\ldots,X_n$.
    - Since X_n is a Markov chain, the relevant rules depend only on the last state X_n so that $I\{T = n\} = (1-d(X_0))\ldots\ (1-d(X_{n-1}))\ d(X_n)$ for some function $d(x)$ defined on S. Such rules are called **stationary**.
    - A stationary rule is, as well, defined by giving partition of $S = S_1 \cup S_2$, where S_1 [S_2] refers to states where the chain is continued [stopped].
  - The goal is to maximize the expected payoff over all stopping rules. Such a rule T^* is called an **optimal stopping time** (or **optimal stopping rule**).
    - Rule T^* is not necessarily unique

4.1 Optimal Stopping of Markov Chains

- *Definition*:
  - The **value** of a state x related to a stopping rule T is the expected payoff assuming the chain starts from x and rule T is used, i.e.

$$v_T(x) = E[f(X_T) \mid X_0 = x]$$

- For any N-vector $u \in R^N$, denote

$$[Pu](x) = \sum_{y \in S} p(x,y)u(y)$$

- For any stationary stopping rule T, we define operator $P_T : R^N \to R^N$,

$$[P_T u](x) = \begin{cases} f(x), & x \in S_2 \\ [Pu](x), & x \in S_1 \end{cases}$$

- *Proposition*:

$$v_T = P_T v_T$$

4.1 Optimal Stopping of Markov Chains

- *Proposition*:
  - For any N-vector $u \in R^N$,

$$\lim_{n \to \infty} (P_T)^n u = v_T$$

- *Proposition*:

$$u = v_T \quad \Leftrightarrow \quad u = P_T u$$

- *Proof*:
  - If $u = v_T$, then we already know that $u = P_T u$.
  - Assume then that $u = P_T u$. Now

$$(P_T)^2 u = P_T(P_T u) = P_T u = u \quad \Rightarrow \quad (P_T)^n u = u \quad \Rightarrow \quad (P_T)^n u \to u$$

  - On the other hand, we know that

$$(P_T)^n u \to v_T$$

  - Thus, $u = v_T$.

4.1 Optimal Stopping of Markov Chains

- *Definition*:
  - The **value** of a state x is the expected payoff assuming the chain starts from x and the optimal rule is used, i.e.

$$v(x) = E[f(X_{T^*}) \mid X_0 = x] = \sup_T E[f(X_T) \mid X_0 = x]$$

- *Proposition*:

$$v = \max\{f, Pv\}$$

- *Idea of the proof*:
  - Consider separately what happens if the chain is stopped or continued at time n
- *Note*:
  - If v is known, then an optimal rule is to stop whenever $v(x) = f(x)$ and continue if $v(x) > f(x)$

4.1 Optimal Stopping of Markov Chains

- *Definition*:
  - Function u is **superharmonic** with respect to P if

$$u \geq Pu$$

- *Note:*
  - Value function $v(x)$ is clearly superharmonic: $v = \max\{f, Pv\} \geq Pv$
- *Proposition*:
  - If u is superharmonic, then for any stopping rule T and any x

$$u(x) \geq E[u(X_T) \mid X_0 = x]$$

- *Proposition*:
  - Value function v is the smallest u such that

$$u \geq \max\{f, Pu\}$$

- *Proof*:

$$u(x) \geq E[u(X_{T^*}) \mid X_0 = x] \geq E[f(X_{T^*}) \mid X_0 = x] = v(x)$$

4.1 Optimal Stopping of Markov Chains

- Define function u_1 as follows:

$$u_1(x) = \begin{cases} f(x), & \text{if } x \text{ is absorbing} \\ \max\{f(y) \mid y \in S\}, & \text{otherwise} \end{cases}$$

- Define function u_n recursively:

$$u_n = \max\{f, Pu_{n-1}\}$$

- *Proposition*:
  - Function u_n is superharmonic and $u_n \geq u_{n+1} \geq f$ for all n
- *Proposition*:

$$\lim_{n\to\infty} u_n = v$$

4.2 Optimal Stopping with Cost

- Suppose P is the transition matrix for a discrete-time Markov chain X_n with a finite state space S
  - As before, let $f(x)$ denote the payoff function and $f^* = \max_x f(x)$
  - Furthermore, associate with each state a **cost** $g(x)$ that must be paid to continue the chain
  - The value function $v(x)$ is thus defined by

$$v(x) = \sup_T E[f(X_T) - \sum_{j=0}^{T-1} g(X_j) \mid X_0 = x]$$

  - And it satisfies:

$$v = \max\{f, Pv - g\}$$

  - In fact, v is the smallest u such that

$$u \geq \max\{u, Pu - g\}$$

  - Algorithm:

$$u_1(x) = \begin{cases} f(x), & \text{if } x \text{ absorbing} \\ f^*, & \text{otherwise} \end{cases} \qquad u_n = \max\{f, Pu_{n-1} - g\} \to v$$

4.3 Optimal Stopping with Discounting

- Suppose P is the transition matrix for a discrete-time Markov chain X_n with a finite state space S
  - As before, let $f(x)$ denote the payoff function and $f^* = \max_x f(x)$
  - Assume now that the value is **discounted** by a factor $0 < \alpha < 1$
  - The value function $v(x)$ is thus defined by

$$v(x) = \sup_T E[\alpha^T f(X_T) \mid X_0 = x]$$

  - And it satisfies:

$$v = \max\{f, \alpha P v\}$$

  - In fact, v is the smallest u such that

$$u \geq \max\{f, \alpha P u\}$$

  - Algorithm:

$$u_1(x) = \begin{cases} f(x), & \text{if } x \text{ absorbing} \\ f^*, & \text{otherwise} \end{cases} \qquad u_n = \max\{f, \alpha P u_{n-1}\} \to v$$

The End of Chapter 4