

Today

- Stopping time problems
- Medical application of stopping time problems

Example: Stopping Time Problems

The following stopping time problem was proposed in 1875

Problem: An employer seeks to hire someone to fill a position. There are N candidates for the position. The candidates are interviewed sequentially. Upon completion of each interview the employer decides whether to offer the job to the current candidate. If no offer is made the individual seeks employment elsewhere and is no longer available for hire.

Stopping Time Problems

Stopping time problems arise when there is a recurring decision to “quit” a process

Basic idea: Each decision epoch the decision maker faces the decision to continue a decision process or quit a receive a one time reward

Example: Parking spot selection



Stopping Time Problems

Epochs: $t = 1, \dots, N$

States: $S = S' \cup \{\Delta\}$, where Δ is an absorbing state

Actions: $A_s = \begin{cases} \{C, Q\}, & s \in S' \\ \{\emptyset\}, & s = \Delta \end{cases}$

Reward if you **C**ontinue

Rewards: $r_t(s, a) = \begin{cases} f_t(s), & s \in S', a = C \\ g_t(s), & s \in S', a = Q, \quad r_N(s) = g_N(s) \\ 0, & s = \Delta \end{cases}$

Reward if you **Q**uit

Transition Probabilities

$$p_t(j|s, a) = \begin{cases} p_t(j|s), & s \in S', j \in S', a = C \\ 1, & s \in S', j = \Delta, a = Q \text{ or } s = j = \Delta \\ 0, & \text{otherwise} \end{cases}$$

Stopping Time Problems

Optimality equations:


$$v_t(s) = \max \left\{ g_t(s), f_t(s, a) + \lambda \sum_{j \in S} p(j|s, C) v_{t+1}(j) \right\}, \forall s \in S$$

$$v_N(s) = g_N(s), \forall s \in S$$

Example: Statin Treatment

Some medical treatment decisions can be viewed as a “stopping time” problem:

- Statins lower your risk of heart attack and stroke
- Treatment has side effects and cost
- Patients receive annual cholesterol tests and make a decision to:
 - initiate statins
 - defer the decision for a year



An advertisement for Crestor (rosuvastatin calcium). On the left is a navigation menu with links: "About CRESTOR", "About cholesterol", "Diet", "Exercise", "Tools for success", and "Important safety information". The main text reads: "Down with the bad cholesterol. CRESTOR® 10 mg, along with diet, can lower bad cholesterol by up to 52% (vs 7% placebo). It can also raise your good cholesterol by up to 14% (vs 3% placebo). Your results may vary." Below this is the phrase "Up with the good." and a button that says "Learn More About CRESTOR®".



Example: Treatment Initiation MDP

Optimality equations:

$$v_t(s) = \max_{a \in \{Q, C\}} \left\{ g_t(s), r(s, C) + \sum_{j \in S} p(j|s, C) v_{t+1}(j) \right\}, \forall s \in S$$

$$v_N(s) = g_N(s), \forall s \in S$$

States are based on patient health status

Action C represents decision to defer statin initiation, Q denotes decision to start treatment

Reward $g_t(s)$ is the expected survival for a patient that has initiated treatment

Absorbing state Δ represents health event that can be avoided by treatment or death from other causes

Markov Decision Process Model

Stages:

- Time horizon: Ages 40-100
- Annual decision epochs


Decision:

- **Initiate** (Q) or **delay** (C) statin treatment

States:

- Metabolic: Total cholesterol and HDL (each can be L, M, H, V)
- Demographic: Gender, Race, BMI, smoking status, medical history

State Transition Diagram



There are various types of reward functions used in health studies. The simplest definition for this problem is:

- $f_t(s_t)$ is the time between decision epochs (e.g. 1 year)
- $g_t(s_t)$ is the expected future life years adjusted for quality of life on medication


Transition probabilities between metabolic states:

- Electronic medical record data for total cholesterol (bad cholesterol) and HDL (good cholesterol) levels for many patients

Transition probabilities from healthy states to complication state


- Published cardiovascular risk models that estimate the probability of heart attack or stroke in the next year

TP Example: Total Cholesterol



Optimal Policy for Diabetic Patients

Bad cholesterol/Good cholesterol



Learn more about these kinds of problems from the following papers in the background reading section on ctools:

- Kurt, M., Denton, B.T., Schaefer, A., Shah, N., Smith, S., “The Structure of Optimal Statin Initiation Policies for Patients with Type 2 Diabetes”, IIE Transactions 1, 49-65, 2011
- Alagoz, O., Maillart, L.M., Schaefer, A.J., Roberts, M.S., “The Optimal Timing of Living-Donor Liver Transplantation,” Management Science, 50(1), 1420-1430

Optimality equations:

$$v_t(s) = \max \left\{ g_t(s), f_t(s, a) + \lambda \sum_{j \in S} p(j|s, C) v_{t+1}(j) \right\}, \forall s \in S$$
$$v_N(s) = g_N(s), \forall s \in S$$

Some stopping time problems have a **control limit** policy:

- Given certain conditions on the reward function and transition probabilities there may exist some ordering of states, S , and control limit, s^* , such that

$$a = \begin{cases} C, & \text{if } s \leq s^* \\ Q, & \text{otherwise} \end{cases}$$

Policies with a simple structure like control limit policies are:

- Easier for decision makers to understand
- Easier to implement
- Easier to solve the associated MDPs

General structure of a **control limit policy**

$$a_t(s_t) = \begin{cases} a_1, & \text{if } s < s^* \\ a_2, & \text{if } s \geq s^* \end{cases}$$


where a_1 and a_2 are alternative actions and s^* is a control limit.

Question: What conditions guarantee the existence of a control limit policy?


Examples of Monotone Policies

Definition: Control limit policies are examples of **monotone** policies. A policy is **monotone** if the *decision rule* at each stage is **nonincreasing** or **nondecreasing** with respect to a chosen ordering of the system states.



nondecreasing
monotone policy
with 2 actions




nondecreasing
monotone policy



More Examples



Q. Which type is this?



Next class we will:

- Define some useful concepts for analyzing MDPs
- Prove theorems about sufficient conditions for a monotone policy

Monotonicity: Optimal Policy

The following provides conditions under which the optimal policy is monotone.

Theorem (4.7.4 Puterman) Suppose for $t = 1, \dots, N - 1$

1. $r_t(s, a)$ is nondecreasing in s for all $a \in A$.
2. $q_t(k|s, a)$ is nondecreasing in s for all $k \in S, a \in A$.
3. $r_t(s, a)$ is **superadditive (subadditive)** on $S \times A$.
4. $q_t(k|s, a)$ is **superadditive (subadditive)** on $S \times A, \forall k$
5. $R_N(s)$ is nondecreasing in s .

Then there exist optimal decision rules, $d_t^*(s)$, which are nondecreasing (nonincreasing) in s for $t = 1, \dots, N - 1$.

Proof: See Puterman, p 107.

Definition: A function is **superadditive** if for $x^+ \geq x^-$, where $x^+, x^- \in X$ and $y^+ \geq y^-$, where $y^+, y^- \in Y$,

$$g(x^+, y^+) + g(x^-, y^-) \geq g(x^+, y^-) + g(x^-, y^+)$$

If the reverse inequality holds then $g(x, y)$ is **subadditive**.

Definition: A Markov chain has the IFR property if there is an ordering of states, $S \equiv \{1, 2, \dots, n\}$, such that

$$q_t(k|s, a) = \sum_{j=k}^n p_t(j|s, a)$$

is nondecreasing in s for all k and a .