

The Optimal Stopping Problem for Multi-armed Bandit Processes

Yuji Yoshida

吉田 祐治

(**Kyushu University**)

§ 0. Abstract. This paper deals with the optimal stopping problem for d -armed bandit processes. Under the assumption of independence of arms this paper shows that not only optimal strategies but also optimal stopping times are expressed by the dynamic allocation indices for each arm. Therefore we see that in order to solve this problem it is sufficient to calculate the dynamic allocation indices for each arm. On the other hand this paper reduces this problem to solve d independent one-parameter optimal stopping problems for each arm and shows that the optimal stopping time for the original problem is expressed explicitly as the sum of d smallest optimal stopping times for one-parameter optimal stopping problems.

Moreover this paper gives a necessary and sufficient condition for the finiteness of the smallest optimal stopping time of the original problem. This condition results in the finiteness of the smallest optimal stopping times of d independent one-parameter stopping problems for each arm.

In Markov case this paper shows that the optimal stopping region is equal to Cartesian product of the optimal stopping regions for each arm and also investigates numerical calculation of optimal strategies and optimal stopping times.

§ 1. Preparation.

§ 1.1. Optimal stopping problem.

$\mathbb{N} = \{0, 1, 2, \dots\}$: time space.

$X = (X_t, \mathcal{F}_t, \mathbb{P})_{t \in \mathbb{N}}$: Markov chain with state space E .

\mathcal{M} : family of all stopping times ($\{\tau = t\} \in \mathcal{F}_t$ for all $t \in \mathbb{N}$, $\mathbb{P}(\tau < \infty) = 1$).

β : discounted rate ($0 < \beta < 1$).

f : bounded measurable function on E .

c : constant function on E .

$Z = (f(X_t))_{t \in \mathbb{N}}$: reward process.

$\mathbb{E}^x \left[\sum_{t=0}^{\tau-1} \beta^t f(X_t) + \beta^\tau c \right]$: expected value of reward process starting from initial state x .

Optimal stopping problem (\mathcal{O}) :

$$\left\{ \begin{array}{l} \text{Find optimal stopping times } \sigma_0 \in \mathcal{M} : \\ \mathbb{E}^x \left[\sum_{t=0}^{\sigma_0-1} \beta^t f(X_t) + \beta^{\sigma_0} c \right] = \sup_{\tau \in \mathcal{M}} \mathbb{E}^x \left[\sum_{t=0}^{\tau-1} \beta^t f(X_t) + \beta^\tau c \right] (= U(x)) \text{ for } x \in E. \end{array} \right.$$

$\sigma_0 = \inf \{ t \in \mathbb{N} : U(X_t) = c \}$: smallest optimal stopping time.

§ 1.2. Dynamic allocation index.

$$v(x) = \sup_{\substack{\tau \in \mathcal{M} \\ \tau \geq 1}} \frac{\mathbb{E}^x \left[\sum_{r=0}^{\tau-1} \beta^r f(X_r) \right]}{\mathbb{E}^x \left[\sum_{r=0}^{\tau-1} \beta^r \right]} \quad (x \in E) : \text{dynamic allocation index}$$

$\tau = \inf \{ t \geq 1 : v(X_t) \leq v(X_0) \}$: optimal stopping time.

§ 2. Notations and results.

§ 2.1. Optimal stopping problem for multi-armed bandit processes.

d : number of arms (positive integer).

β : discounted rate ($0 < \beta < 1$).

$X^i = (X_t^i, \mathfrak{F}_t^i, \mathbb{P}^i)_{t \in \mathbb{N}}$: mutually independent Markov chains with state space E^i ($i = 1, \dots, d$).

f^i : bounded measurable function on E^i ($i = 1, \dots, d$).

c : constant function on E .

$Z^i = (f^i(X_t^i))_{t \in \mathbb{N}}$: machine with arm i (= reward process i) ($i = 1, \dots, d$).

$X = (X_s)_{s \in \mathbb{T}} = (X_{s1}^1, \dots, X_{sd}^d)_{s = (s^1, \dots, s^d) \in \mathbb{T}}$: d -parameter process with state space E .

$$\left\{ \begin{array}{l} \mathbb{T} = \mathbb{N}^d : \text{time space} \quad E = \prod_{i=1}^d E^i : \text{state space} \\ \mathbb{P} = \prod_{i=1}^d \mathbb{P}^i : \text{probability} \quad \mathfrak{F}_t = \bigotimes_{i=1}^d \mathfrak{F}_{t_i}^i : \sigma\text{-field } (t = (t_1, \dots, t_d)) \end{array} \right.$$

strategy $\pi = (\pi(t))_{t \in \mathbb{N}} = ((\pi^1(t), \dots, \pi^d(t)))_{t \in \mathbb{N}}$

$\pi(t) : \Omega \rightarrow \mathbb{T}$ satisfying (i), (ii) and (iii):

(i) $\pi(0) = (0, \dots, 0)$.

(ii) For all $t \in \mathbb{N}$ it holds that

$$\pi(t+1) = \pi(t) + (0, \dots, \overset{\downarrow 1}{0, 1, 0, \dots}, 0) \quad \text{for some } i = 1, \dots, d.$$

(iii) $(\pi(t) = r) \in \mathfrak{F}_t$ for all $t \in \mathbb{N}$ and all $r \in \mathbb{T}$.

$\pi^i(t)$: number of pulls of arm i up to time t (= time of machine i at time t).

\mathcal{A} = (all strategies π).

$\mathfrak{M}^\pi = \{ \tau : \Omega \rightarrow \mathbb{N} \text{ s.t. } \{ \tau = t \} \cap \{ \pi(t) = r \} \in \mathcal{F}_r \text{ for } t \in \mathbb{N} \text{ and } r \in \mathbb{T} \}$

: family of all stopping times along $\pi \in \mathcal{I}$.

$$V^{\pi, \tau}(x) = \mathbb{E}^x \left[\sum_{t=0}^{\tau-1} \sum_{i=1}^d \beta^t r^i(x_{\pi^i(t)}) (\pi^i(t+1) - \pi^i(t)) + \beta^\tau c \right] \quad (\pi \in \mathcal{I} \text{ and } \tau \in \mathfrak{M}^\pi)$$

: expected value of reward process.

$$V^{*,*}(x) = \sup_{\pi \in \mathcal{I}} \sup_{\tau \in \mathfrak{M}^\pi} V^{\pi, \tau}(x) \quad (x = (x^1, \dots, x^d) \in E) : \text{optimal value.}$$

Optimal stopping problem for d-armed bandit processes ($\mathcal{O/B}$):

$$\left\{ \begin{array}{l} \text{Find strategies } \pi^* \in \mathcal{I} \text{ and stopping times } \tau^* \in \mathfrak{M}^{\pi^*} \\ V^{\pi^*, \tau^*}(x) = \sup_{\pi \in \mathcal{I}} \sup_{\tau \in \mathfrak{M}^\pi} V^{\pi, \tau}(x) (= V^{*,*}(x)) \text{ for } x \in E \end{array} \right.$$

Theorem 1 (Bellman equation)

$$V^{*,*}(x) = \max \left\{ c, \max_{1 \leq i \leq d} \left\{ r^i(x^i) + \mathbb{E}^x \left[\beta V^{*,*}(X_1^i) \right] \right\} \right\}.$$

§ 2.2. Optimal strategies and Optimal stopping times.

$v^i(x^i)$ ($i = 1, \dots, d$): dynamic allocation index with arm i .

$$v(x) = \max_{1 \leq i \leq d} v^i(x^i) \quad (x = (x^1, \dots, x^d) \in E) : \text{maximum index.}$$

Theorem 2

(i) $\pi^* (\in \mathcal{I})$: index strategy i.e.

for each $t \in \mathbb{N}$ π^* satisfies

$$v(\mathbf{X}_{\pi^*(t)}) = v^i(\mathbf{X}_{\pi^*(t)}^i) \text{ on } (\pi(t+1) = \pi(t) + (0, \dots, 0, \overset{i}{1}, 0, \dots, 0))$$

for some $i = 1, \dots, d$.

(ii) $\tau^{\pi^*} = \inf \{ t \in \mathbb{N} : v(\mathbf{X}_{\pi^*(t)}) \leq (1 - \beta) \cdot c \}$.

(iii) $\mathbb{P}(\tau^{\pi^*} < \infty) = 1$

$\Rightarrow \pi^*$ is optimal strategy and τ^{π^*} is optimal stopping time for $(\mathcal{O}, \mathcal{B})$.

\mathfrak{M}^i : family of all \mathcal{F}^i - adapted stopping times $(\tau \leq t) \in \mathcal{F}_t^i$ for all $t \in \mathbb{N}$, $\mathbb{P}(\tau < \infty) = 1$.

Optimal stopping problem with arm i (\mathcal{O}^i):

$$\left\{ \begin{array}{l} \text{Find optimal stopping times } \sigma_0^i \in \mathfrak{M}^i : \\ \mathbb{E}^{\mathbf{X}^i} \left[\sum_{t=0}^{\sigma_0^i-1} \beta^t f(\mathbf{X}_t^i) + \beta^{\sigma_0^i} c \right] = \sup_{\tau \in \mathfrak{M}^i} \mathbb{E}^{\mathbf{X}^i} \left[\sum_{t=0}^{\tau-1} \beta^t f(\mathbf{X}_t^i) + \beta^\tau c \right] (= U^i(\mathbf{X}^i)) \text{ for } \mathbf{X}^i \in E^i \end{array} \right.$$

$\sigma_0^i = \inf \{ t \in \mathbb{N} : U^i(\mathbf{X}_t^i) = c \}$ ($i = 1, \dots, d$): its optimal stopping time.

Theorem 3

$\pi^* \in \mathcal{H}$: index strategy.

\Rightarrow

(i) $\sigma_0^i = \inf \{ t \in \mathbb{N} : v^i(\mathbf{X}_t^i) \leq (1 - \beta) \cdot c \} = \pi^{*i}(\tau^{\pi^*})$ ($i = 1, \dots, d$). ... machine i

(ii) $\tau^{\pi^*} = \sum_{i=1}^d \sigma_0^i = \inf \{ t \in \mathbb{N} : V^{*,*}(\mathbf{X}_{\pi^*(t)}) = c \}$ bandit process

(iii) $\mathbb{P}[\tau^{\pi^*} < \infty] = \prod_{i=1}^d \mathbb{P}[\sigma_0^i < \infty]$.