Incentivizing Hidden Types in Secretary Problem

Longjian Li[1] Alexis Akira Toda[2]

[1]Peking University

[2]University of California San Diego

Seminar @UCSD
October 7, 2022

Classical secretary problem

- Administrator sequentially interviews job applicants $1, \ldots, N$ in random order
- Can rank applicants already interviewed from best to worst
- Must accept or reject applicant immediately after interview, with no recall
- What is optimal stopping rule to maximize probability of hiring the best?

Classical secretary problem

- Administrator sequentially interviews job applicants $1, \ldots, N$ in random order
- Can rank applicants already interviewed from best to worst
- Must accept or reject applicant immediately after interview, with no recall
- What is optimal stopping rule to maximize probability of hiring the best?
- Applications I have in mind:
  - Film director seeks to identify best fit actor
  - Department seeks to hire best junior candidate
  - Racquet manufacturer seeks to sponsor next Rafael Nadal

Solution to classical secretary problem

- Define threshold

$$n^* = \min\left\{ n : \sum_{k=n}^{N-1} \frac{1}{k} \leq 1 \right\}.$$

- Reject first $n^* - 1$ applicants
- Accept next applicant if best among those already interviewed, otherwise reject
- As $N \to \infty$, we can show

$$\frac{n^*}{N} \to \frac{1}{\mathrm{e}} \approx 0.37$$
$$\Pr(\text{success}) \to \frac{1}{\mathrm{e}} \approx 0.37$$

Question

- First $n^* - 1$ applicants always rejected, so no incentive to show up for interviews
- If applicants don't show up, administrator can't learn applicants' abilities
- What is optimal strategy of administrator if applicants incur cost $c \in [0, 1)$ (relative to job value) to complete interview and must be incentivized to show up?

This paper

- Prove existence of unique full learning equilibrium
  - Administrator can tell whether current applicant is best among those already invited for interviews
- Prove optimality of full learning equilibrium
  - Among all equilibria, full learning equilibrium achieves maximum success probability
- Characterize asymptotic behavior as $N \to \infty$
  - Success probability π_N^* exhibits power law decay N^{-c}

Agents

- An administrator
- $N \geq 2$ job applicants, invited for interviews in order $1, 2, \ldots, N$
- Applicant n has ability $\theta_n > 0$

Actions

- When invited for interview, applicant chooses action $a = 0$ (decline interview) or $a = 1$ (complete interview)
- Interview reveals output $y = a\theta$, where θ: ability
- Immediately after interview n, administrator must accept or reject applicant n based only on history of observed outputs $\{y_1, \ldots, y_n\}$
- Game ends if applicant accepted; move to next applicant if rejected; no recall

Payoffs

- Applicant:
  - Job value normalized to 1
  - Completing interview costs $c \in [0, 1)$
- Administrator: if accept applicant with ability θ, then payoff is

$$\begin{cases} 1 & \text{if } \theta = \max_{1 \leq n \leq N} \theta_n, \\ 0 & \text{otherwise.} \end{cases}$$

- All agents risk-neutral

Information

- Abilities $\{\theta_n\}_{n=1}^N$ realized before game begins but private information
- Administrator believes rank orders of $\{\theta_n\}_{n=1}^N$ have no ties and equally likely with probability $1/N!$
- When administrator invites applicant n, presents past outputs $\{y_1, \ldots, y_{n-1}\}$
- Applicant chooses action $a_n \in \{0, 1\}$ and output $y_n = a_n\theta_n$ observed
- After game ends, $\{\theta_n\}_{n=1}^N$ becomes public information and payoffs realized

Strategies

- Let $H_n = \mathbb{R}_+^n$ be set of outputs of first n applicants ($H_0 = \emptyset$)
- Applicant n's strategy is a function $s_n : H_{n-1} \times (0, \infty) \to \{0, 1\}$
  - $s_n(y_1, \ldots, y_{n-1}, \theta) = 1$ $(= 0)$ means applicant n with ability θ completes (declines) interview given past outputs $(y_1, \ldots, y_{n-1})$
- Administrator's (mixed) strategy is a collection of functions $\sigma = \{\sigma_n\}_{n=1}^N$ with $\sigma_n : H_n \to [0, 1]$
  - $p = \sigma_n(y_1, \ldots, y_n)$ is probability administrator accepts applicant n given outputs $(y_1, \ldots, y_n)$
  - commitment power, so choose σ once and for all
- Nash equilibrium is strategy profile $(\sigma^*, s_1^*, \ldots, s_N^*)$ that is mutually best response

Full learning equilibrium

- Focus on *full learning equilibrium*
- We say equilibrium is full learning if for any equilibrium path and n until game ends, we have

$$\max_{1 \le k \le n} \theta_k = \max_{1 \le k \le n} y_k$$

- This condition allows administrator to tell if current applicant is best among those already interviewed

Lemma

Let $(\sigma^*, s_1^*, \ldots, s_N^*)$ *be a full learning equilibrium. Then* $\theta_1 = y_1$ *and*

$$\theta_n \begin{cases} = \max_{1\le k\le n} \theta_k & \text{if } y_n > \max_{1\le k\le n-1} y_k, \\ < \max_{1\le k\le n} \theta_k & \text{if } y_n \le \max_{1\le k\le n-1} y_k \end{cases}$$

for $n \geq 2$ *until the game ends.*

Lemma

Let $(\sigma^*, s_1^*, \ldots, s_N^*)$ *be a full learning equilibrium. Then* $\theta_1 = y_1$ *and*

$$\theta_n \begin{cases} = \max_{1 \le k \le n} \theta_k & \text{if } y_n > \max_{1 \le k \le n-1} y_k, \\ < \max_{1 \le k \le n} \theta_k & \text{if } y_n \le \max_{1 \le k \le n-1} y_k \end{cases}$$

for $n \geq 2$ *until the game ends.*

Proof.

- If $y_n > \max_{1 \le k \le n-1} y_k$, then $0 < y_n = a_n \theta_n$ so $a_n = 1$ and $\theta_n = y_n$
- Hence $\theta_n = y_n = \max_{1 \le k \le n} y_k = \max_{1 \le k \le n} \theta_k$
- If $y_n \le \max_{1 \le k \le n-1} y_k$, then

$$\theta_n \le \max_{1 \le k \le n} \theta_k = \max_{1 \le k \le n} y_k = \max_{1 \le k \le n-1} y_k = \max_{1 \le k \le n-1} \theta_k,$$

and inequality strict because no ties □

Partial characterization of equilibrium strategy

Lemma

Let $(\sigma^*, s_1^*, \ldots, s_N^*)$ *be a full learning equilibrium. Then*

$$\sigma_n^*(y_1, \ldots, y_n) \begin{cases} \geq c & \text{if } y_n > \max_{1 \leq k \leq n-1} y_k, \\ = 0 & \text{if } y_n \leq \max_{1 \leq k \leq n-1} y_k, \end{cases}$$

$$s_n^*(y_1, \ldots, y_{n-1}, \theta) = \begin{cases} 1 & \text{if } \theta > \max_{1 \leq k \leq n-1} y_k, \\ 0 & \text{if } \theta \leq \max_{1 \leq k \leq n-1} y_k. \end{cases}$$

Partial characterization of equilibrium strategy

Lemma

Let $(\sigma^*, s_1^*, \ldots, s_N^*)$ be a full learning equilibrium. Then

$$\sigma_n^*(y_1, \ldots, y_n) \begin{cases} \geq c & \text{if } y_n > \max_{1 \leq k \leq n-1} y_k, \\ = 0 & \text{if } y_n \leq \max_{1 \leq k \leq n-1} y_k, \end{cases}$$

$$s_n^*(y_1, \ldots, y_{n-1}, \theta) = \begin{cases} 1 & \text{if } \theta > \max_{1 \leq k \leq n-1} y_k, \\ 0 & \text{if } \theta \leq \max_{1 \leq k \leq n-1} y_k. \end{cases}$$

Idea:

- Administrator promises acceptance probability c if applicant current best so as to incentivize completing interview
- Then current applicant completes interview if current best, otherwise declines

Dynamic programming

- State is n and $x \in X = \{0, 1\}$, where
  - $x = 1$: applicant is current best
  - $x = 0$: applicant is not current best
- By random order, we have $\Pr(x' = 1) = \frac{1}{n+1}$ independent of x
- Let $V_n(x)$ be value function; then Bellman equation is

$$V_n(0) = \frac{1}{n+1} V_{n+1}(1) + \frac{n}{n+1} V_{n+1}(0)$$

- If $x = 1$, need to promise acceptance probability $p \geq c$ to incentivize applicant to complete interview

Dynamic programming

- If accept, payoff is

$$\begin{aligned}
&\Pr\left(\theta_n = \max_{1\le k\le N} \theta_k \mid x = 1\right)\\
&= \Pr(n \text{ is best among all} \mid n \text{ is best among first } n)\\
&= \Pr(n \text{ is best among all and first } n) / \Pr(n \text{ is best among first } n)\\
&= \Pr(n \text{ is best among all}) / \Pr(n \text{ is best among first } n)\\
&= (1/N)/(1/n) = \frac{n}{N}
\end{aligned}$$

- Hence Bellman equation is

$$V_n(1) = \max_{c\le p\le 1}\left\{p\frac{n}{N} + (1-p)\left(\frac{1}{n+1}V_{n+1}(1) + \frac{n}{n+1}V_{n+1}(0)\right)\right\}$$

Dynamic programming

Proposition

The value functions in a full learning equilibrium satisfy $V_N(0) = 0$, $V_N(1) = 1$, *and*

$$V_n(0) = \frac{1}{n+1} V_{n+1}(1) + \frac{n}{n+1} V_{n+1}(0),$$

$$V_n(1) = \max_{c \le p \le 1} \left\{ p\frac{n}{N} + (1-p)V_n(0) \right\}$$

$$= \max \left\{ c\frac{n}{N} + (1-c)V_n(0), \frac{n}{N} \right\} > 0$$

- Define normalized value function $v_n(x) := V_n(x)/n$
- Dividing Bellman equations by n, we get

Normalized value functions

Proposition

The normalized value $v_n(x) = V_n(x)/n$ satisfies $v_N(0) = 0$, $v_N(1) = 1/N$, and

$$v_n(0) = \frac{1}{n} v_{n+1}(1) + v_{n+1}(0),$$
$$v_n(1) = \max\{c/N + (1-c)v_n(0), 1/N\}.$$

Furthermore, $v_n(0)$ is strictly decreasing in n and $v_n(1)$ is decreasing in n.

- Strict monotonicity of $v_n(0)$ implies that there exists threshold n^* such that
  - Accept current best applicant n with probability c if $n < n^*$
  - Accept current best applicant n with probability 1 if $n \geq n^*$

Existence of full learning equilibrium

Theorem

For all $N \geq 2$ and $c \in [0,1)$, there exists a unique full learning equilibrium, which can be constructed as follows:

1. *Define* $n^* = \min\left\{n : \sum_{k=n}^{N-1} \frac{1}{k} \leq 1\right\}$.
2. *Define* $\sigma_n^* : H_n \to [0,1]$ *by*

$$\sigma_n^*(y_1, \ldots, y_n) = \begin{cases} 1 & \text{if } n \geq n^* \text{ and } 0 < y_n = \max_{1 \leq k \leq n} y_k, \\ c & \text{if } n < n^* \text{ and } 0 < y_n = \max_{1 \leq k \leq n} y_k, \\ 0 & \text{otherwise.} \end{cases}$$

3. *Define* $s_n^* : H_{n-1} \times (0, \infty) \to \{0,1\}$ *by*

$$s_n^*(y_1, \ldots, y_{n-1}, \theta) = \begin{cases} 1 & \text{if } \theta > \max_{1 \leq k \leq n-1} y_k, \\ 0 & \text{if } \theta \leq \max_{1 \leq k \leq n-1} y_k. \end{cases}$$

Optimality of full learning equilibrium

- This game has many equilibria (e.g., ignore first k candidates and then learn)
- Which equilibrium is best?

Optimality of full learning equilibrium

- This game has many equilibria (e.g., ignore first k candidates and then learn)
- Which equilibrium is best?

Theorem

The full learning equilibrium is optimal in the sense that the success probability is the highest among all equilibria.

Proof idea

- Proof is difficult because there are many ways to deviate (learn or not learn)
- Let $P_{n,N}(y_1, \ldots, y_n)$ be success probability in any equilibrium conditional on interviewing first n applicant and full learning
- For $n+1$, possible deviations are (i) learn and accept with probability $p \in [c, 1]$ conditional on current best, or (ii) not learn and accept with probability $p \in [0, 1]$
- Use induction on $j = N - n$ (number of remaining applicants) to bound $P_{n,N}(y_1, \ldots, y_n)$ from above by continuation value of full learning equilibrium
- Then full learning equilibrium is optimal because $P_{1,N} \leq V_1(1)$ by induction

Asymptotic behavior: threshold n^*

Proposition

The threshold $n_N^* = \min\left\{n : \sum_{k=n}^{N-1} \frac{1}{k} \leq 1\right\}$ *satisfies*

$$\frac{N}{\mathrm{e}} \leq n_N^* \leq \frac{N-1}{\mathrm{e}} + 2.$$

In particular, $\lim_{N\to\infty} n_N^*/N = 1/\mathrm{e} = 0.367\ldots.$

Proof

- Let $t = n_N^*$
- By definition, we have

$$1 \geq \sum_{k=t}^{N-1} \frac{1}{k} \geq \int_t^N \frac{1}{x}\,\mathrm{d}x = \log \frac{N}{t} \implies t \geq \frac{N}{\mathrm{e}}$$

- Similarly,

$$1 < \sum_{k=t-1}^{N-1} \frac{1}{k} \leq \int_{t-2}^{N-1} \frac{1}{x}\,\mathrm{d}x = \log \frac{N-1}{t-2} \implies t \leq \frac{N-1}{\mathrm{e}} + 2$$

Asymptotic behavior: success probability $\pi_N^* = V_1(1)$

Theorem

If $c \in [0, 1)$, *then*

$$\lim_{N\to\infty} N^c \pi_N^* = \frac{\mathrm{e}^{c-1}}{\Gamma(2-c)},$$

where Γ *is the gamma function. In particular, if* $c = 0$ *then* $\lim_{N\to\infty} \pi_N^* = 1/\mathrm{e} = 0.367\ldots$.

Asymptotic behavior: success probability $\pi_N^* = V_1(1)$

Theorem

If $c \in [0,1)$, then

$$\lim_{N\to\infty} N^c \pi_N^* = \frac{\mathrm{e}^{c-1}}{\Gamma(2-c)},$$

where Γ is the gamma function. In particular, if $c = 0$ then $\lim_{N\to\infty} \pi_N^* = 1/\mathrm{e} = 0.367\ldots$.

- π_N^* exhibits a power law decay with exponent $-c$
- Proof uses value function iteration, Gauss product formula for gamma function

$$\Gamma(z) = \lim_{n\to\infty} \frac{n^z n!}{z(z+1)\cdots(z+n)},$$

  and definition of Riemann integral (pretty cool)

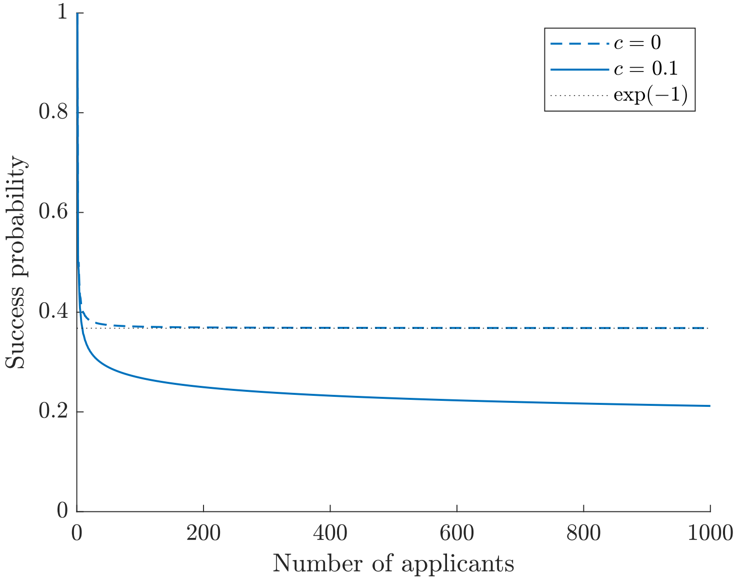
$c = 0$
$c = 0.1$
exp(−1)
Success probability
Number of applicants
0
0.2
0.4
0.6
0.8
1
0
200
400
600
800
1000

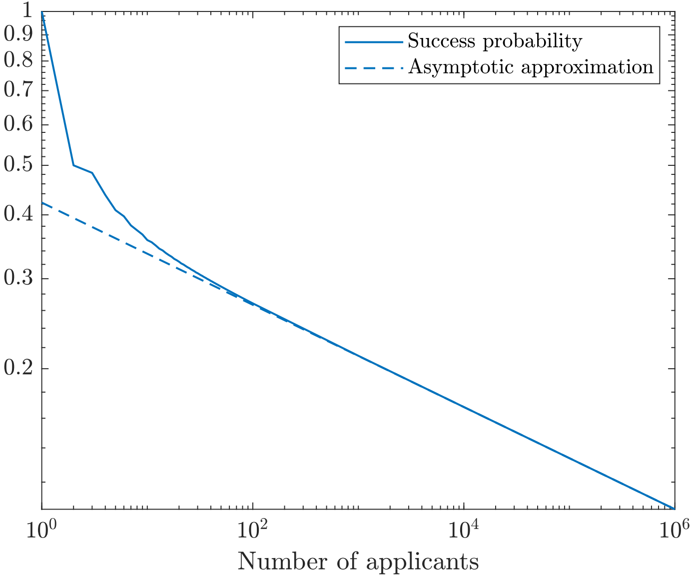
Success probability
Asymptotic approximation
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
10^0
10^2
10^4
10^6
Number of applicants

Concluding remarks

- Extended classical secretary problem with incentives
- Even if applicants incur interview cost c, threshold is same as no cost
- However, for applicants $n < n^*$, administrator accepts with probability c if current best to incentivize completing interview
- Future work: what if administrator observes noisy signal of ability θ_n and can decide interview order?

References