Prophet Secretary Through Blind Strategies

J.Correa [1], **Raimundo Saona**[1], B. Ziliotto [2]

[1]Universidad de Chile

[2]CEREMADE, CNRS, Université Paris Dauphine

Classical settings, Secretary Problem

Classical settings, Secretary Problem

$$\downarrow$$

10

Classical settings, Secretary Problem

$\downarrow$ $\qquad$ $\downarrow$

$\not{1}\not{0}$ $\qquad$ π

Classical settings, Secretary Problem

$\downarrow$	$\downarrow$	$\downarrow$	
98	~~10~~	~~π~~	

Classical settings, Secretary Problem

↓	↓	↓	↓
~~98~~	~~10~~	~~π~~	0.4

Classical settings, Prophet Inequality

F_1

F_2

F_3

F_4

Classical settings, Prophet Inequality

F_1

$\downarrow$

2

F_2

F_3

F_4

Classical settings, Prophet Inequality

   

F_1	F_2	F_3	F_4
$\downarrow$	$\downarrow$		
$\not{2}$	8		

Classical settings, Prophet Inequality

F_1	F_2	F_3	F_4
$\downarrow$	$\downarrow$	$\downarrow$	
$\not{2}$	$\not{8}$	3	

What will we discuss?

The classical **Prophet Inequality**

has connections with online sales
through **posted price mechanism**.

Prophet Secretary = Prophet Inequality with **random arrival**.

Formal dynamics

1. You are given $F_1, \dots, F_n$ distributions over $[0, \infty)$.
2. Following a **uniform random order** σ, you are shown

$$V_{\sigma_1} \sim F_{\sigma_1} .$$

3. If V_{σ_1} was taken, the process ends.
4. If not, you are shown the pair

$$V_{\sigma_2} \sim F_{\sigma_2} .$$

5. If V_{σ_2} was taken, the process ends.
6. ...

Note: $V_1, \dots, V_n$ are independent random variables.

Performance measure

For every instance $(F_1, \ldots, F_n)$, the player can choose a selection algorithm ALG. The performance for this instance will be

$$\frac{\mathbb{E}\,(ALG)}{\mathbb{E}\left(\max_{i\in[n]}\{V_i\}\right)}\,.$$

And a family of algorithms is said to have a perform of c if

$$\inf_{F_1,\ldots,F_n} \frac{\mathbb{E}\,(ALG)}{\mathbb{E}\left(\max_{i\in[n]}\{V_i\}\right)} \geq c\,.$$

In this setting, it is still **unknown the performance** of the optimal algorithm given by dynamic programming.

Slight variation

It is remarkable that there is no result that separates, in terms of **achievable performance**, the following three settings.

1. Prophet Secretary: c_{ProSec}
   $F_1, \dots, F_n$ different distributions and
   σ independent uniform random order.
2. Order Selection: c_{OrdSel}
   $F_1, \dots, F_n$ different distributions and
   σ order chosen by the player.
3. I.I.D. Prophet Inequality: c_{iid}
   $F_1 = \dots = F_n = F$ a fixed distribution for everyone.

$$c_{ProSec} \leq c_{OrdSel} \leq c_{iid}$$

Main results

Algorithms

$1 - 1/e \approx 0.632$	$\leq$	0.635	$\leq$	0.669	$\leq$	$\overline{c}$
Previous algorithms		Azar et al. 2018		THIS PAPER 2019		

Upper bound

$\overline{c}$	$\leq$	0.675	$\leq$	0.732	$\leq$	0.745
		blind		nonadaptive		(IID Case)
		THIS PAPER 2019				Hill & Kertz 1982 Correa et al. 2018

Fixed threshold

Theorem[Ehsani et al. 2018]
A fixed threshold algorithm can achieve a performance of $1 - 1/e$.

Proof(continuous case). Compute τ such that

$$\mathbb{P}(max \leq \tau) = 1/e\,.$$

$ALG_\tau :=$ **pick the first value above** τ.

Note that if $t \leq \tau$,
$\mathbb{P}(ALG_\tau > t) = \mathbb{P}(ALG_\tau > 0) = 1 - 1/e \geq (1 - 1/e)\mathbb{P}(\max > t)$.

Lemma. $\mathbb{P}(\text{pick } V_i | V_i > \tau) \geq 1 - 1/e$.

Fixed threshold (continuation)

If $t > \tau$,

$$
\begin{aligned}
\mathbb{P}(ALG_\tau > t) &= \sum_{i \le n} \mathbb{P}(V_i > t |\text{ pick } V_i)\, \mathbb{P}(\text{pick } V_i) \\
&= \sum_{i \le n} \frac{\mathbb{P}(V_i > t)}{\mathbb{P}(V_i > \tau)}\, \mathbb{P}(\text{pick } V_i) \\
&= \sum_{i = \le n} \mathbb{P}(V_i > t)\, \mathbb{P}(\text{pick } V_i | V_i > \tau) \\
&\ge (1 - 1/e) \sum_{i \le n} \mathbb{P}(V_i > t) \ge (1 - 1/e)\, \mathbb{P}(\max > t)\,.
\end{aligned}
$$

Integrating on t, we get $\mathbb{E}(ALG_\tau) \ge (1 - 1/e)\mathbb{E}(\max)$.

Analysing time-thresholds

A fixed threshold can achieve a performance of $1 - 1/e$,
What if we use a threshold for the first half and then another?

Time thresholds could be analysed further.
Blind strategies is just one way of defining these thresholds.
Given $\alpha : [0,1] \to [0,1]$, define $\tau_1, \ldots, \tau_n$ by

$$\mathbb{P}(\max_{i \in [n]}\{V_1, \ldots, V_n\} \leq \tau_i) = \alpha(i/n)\,.$$

Note: α is **instance-independent** and **quantile-based**.

Simplified definition

Fix $\alpha : [0,1] \to [0,1]$
Given an instance $F_1, \ldots, F_n$, compute $\tau_1, \ldots, \tau_n$ such that

$$\mathbb{P}(\max_{i \in [n]}\{V_1, \ldots, V_n\} \leq \tau_i) = \alpha(i/n)\,.$$

Then, **accept** V_{σ_i} **if it is larger than** τ_i**.**

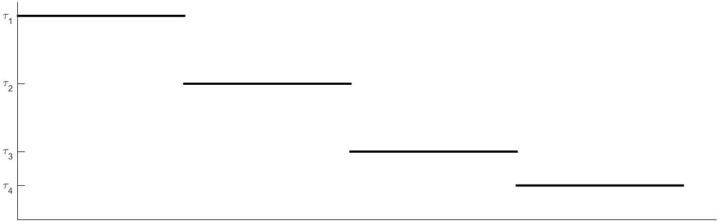

Simplified definition

Fix $\alpha : [0,1] \to [0,1]$
Given an instance $F_1, \ldots, F_n$, compute $\tau_1, \ldots, \tau_n$ such that

$$\mathbb{P}(\max_{i \in [n]}\{V_1, \ldots, V_n\} \leq \tau_i) = \alpha(i/n)\,.$$

Then, **accept V_{σ_i} if it is larger than τ_i.**

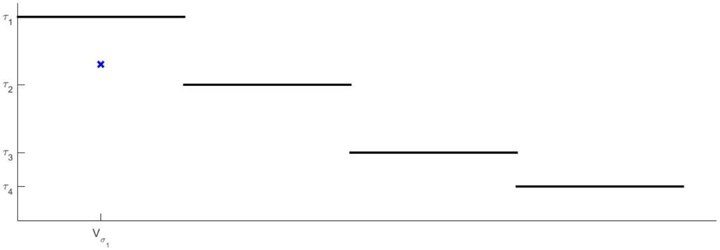

Simplified definition

Fix $\alpha : [0,1] \to [0,1]$
Given an instance $F_1, \ldots, F_n$, compute $\tau_1, \ldots, \tau_n$ such that

$$\mathbb{P}(\max_{i \in [n]}\{V_1, \ldots, V_n\} \leq \tau_i) = \alpha(i/n)\,.$$

Then, **accept V_{σ_i} if it is larger than τ_i.**

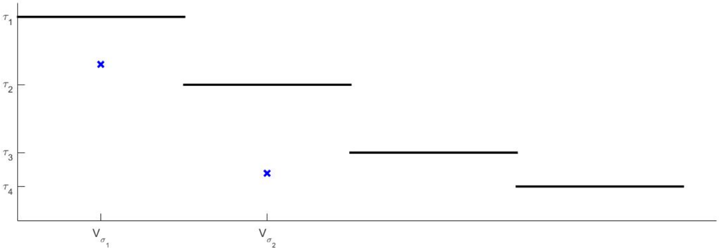

Simplified definition

Fix $\alpha : [0,1] \to [0,1]$
Given an instance $F_1, \ldots, F_n$, compute $\tau_1, \ldots, \tau_n$ such that

$$\mathbb{P}(\max_{i \in [n]}\{V_1, \ldots, V_n\} \leq \tau_i) = \alpha(i/n)\,.$$

Then, **accept** V_{σ_i} **if it is larger than** τ_i**.**

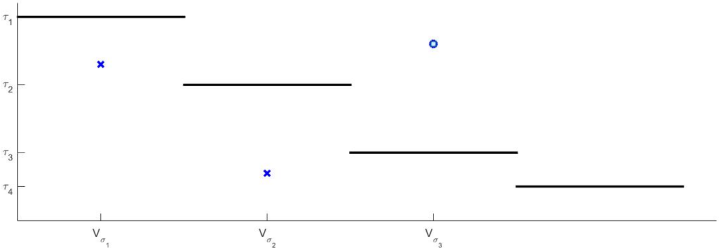

Characteristics

Blind strategies are:

1. **Anonymous**: Do not take the identity of the variable revealed into account, only the value observed.
2. **Nonadaptive or Static**: Do not take previous values or identity of observed variables into account when facing new values.
3. **Quantile based**: Compare values only based on quantiles of the distribution of the maximum.
4. **Instance-independent**: The choice of α does not depend on the instance.

General idea

To prove

$$\mathbb{E}(ALG) \geq \overline{c}\, \mathbb{E}(\max)\,,$$

we prove that, for all $t > 0$,

$$\mathbb{P}(ALG > t) \geq \overline{c}\, \mathbb{P}(\max > t)\,.$$

For this, fix $\tau_1 \geq \ldots \geq \tau_n$ and study the intervals

$$[0, \tau_n), [\tau_n, \tau_{n-1}), \ldots, [\tau_2, \tau_1), [\tau_1, \infty)\,.$$

Interval decomposition

For $t \in [\tau_j, \tau_{j-1})$,

$$\mathbb{P}(ALG > t) \geq c(F_T, \alpha, j)\mathbb{P}(\max > t) .$$

Technical challenge:

How to relate the distribution of the stopping time F_T with the choice of α?

F_T and $F_{\max}$

The key result is that the distribution of the stopping time F_T

1. **Upper bound**: Is maximized in the i.i.d. case.
2. **Lower bound**: Is minimized in the all-but-one not null case.

$$g_\alpha \leq F_T \leq h_\alpha$$

We can **drop the dependence** on F_T and work only with α, so

$$\mathbb{P}(ALG > t) \geq c_\alpha(j)\mathbb{P}(\max > t)\,.$$

Optimizing over the choice of α, we find that there is α^* with good performance **in every interval**, ie: for all j,

$$c_{\alpha^*}(j) \geq \overline{c} \geq 0.669$$

Summing up

Then, for

$$\begin{aligned}\mathbb{P}(ALG > t) &\geq c_{\alpha^*}(j)\,\mathbb{P}(\max > t)\\ &\geq \overline{c}\,\mathbb{P}(\max > t)\\ &\geq 0.669\,\mathbb{P}(\max > t)\,.\end{aligned}$$

Integrating over t, we conclude

$$\mathbb{E}(ALG) \geq 0.669\,\mathbb{E}(\max)$$

An upper bound for static algorithms

No "nonadaptive" algorithm can achieve a better performance than $\sqrt{3}-1 \approx 0.732$. The algorithm can depend on: the value observed, the identity and on time, but not on the history.

Key instance

$$V_1 \equiv \delta = \sqrt{3} - 1$$
$$V_2 = \begin{cases} n & w.p. \quad 1/n \\ 0 & w.p. \quad 1 - 1/n \end{cases}$$
$$V_3 = \ldots = V_n \equiv 0 \,.$$

The algorithm always picks n, if faced, and must decide

faced with δ at time i, how likely do I accept it?

Optimal strategy and performance

For every n and $\delta \in [0,1)$,

1. If you face δ **too early**, probably V_2 appears after (we do not remember if it already appeared), so we should **not pick it**.
2. If you face δ **late**, because probably you already faced V_2, just **pick it**.

This defines the optimal algorithm ALG^* and

$$\lim_{n\to\infty} \frac{\mathbb{E}(ALG^*)}{\mathbb{E}(max)} = \frac{1+\delta^2/2}{1+\delta} = \sqrt{3}-1 \approx 0.732\,.$$

Summary

Algorithms

$1 - 1/e \approx 0.632$	$\leq$	0.635	$\leq$	0.669	$\leq$	$\bar{c}$
Previous algorithms		Azar et al. 2018		THIS PAPER 2019		

Upper bound

$\bar{c}$	$\leq$	0.675	$\leq$	0.732	$\leq$	0.745
		blind		nonadaptive		(IID Case)
		THIS PAPER 2019				Hill & Kertz 1982 Correa et al. 2018

Future directions

Three different(?) settings

1. Prophet Secretary (random order)
    - How good are nonadaptive algorithms?
    - Is the optimal performance worse than the i.i.d. case?
2. Order Selection (free-order)
    - Complexity?
    - How to compute the best order?
    - Can we get the i.i.d. performance?
    - Can nonadaptive algorithms be as good as the optimal one?
3. I.I.D. Prophet Inequality (all have the same distribution)