Prophet Inequalities and Prophet Secretary

Andrew Chen Nathan Finkle

December 2021

1 Introduction

In this paper, we hope to outline different approximation algorithms for the **Prophet Inequalities** and **Prophet Secretary** problems and suggest some extensions from the **Prophet Inequalities** algorithms to the **Prophet Secretary** problem.

1.1 Prophet Inequalities & Prophet Secretary Definitions

Our setting has 1 item and n buyers with known (possibly distinct) distributions $\{D_1, \ldots, D_n\}$. When a buyer i arrives, they tell you a price X_i drawn from D_i, and you must choose whether to sell your one item to them at X_i or permanently pass on selling to i. Our goal is to maximize revenue, and we will compare our algorithms' revenue to the optimal revenue $\text{OPT} = \mathbb{E}\left[\max_{i=1,\ldots,n} X_i\right]$.[1]

In **Prophet Inequalities**, the buyers arrive in an adversarial order known in advance to the algorithm. That is, at every step $i \in [n]$, price $X_i \sim D_i$ is offered.

In **I.I.D. Prophet Inequalities**, the distributions are all the same. That is, at every step $i \in [n]$, price $X_i \sim D$ is offered.

In **Prophet Secretary**, the order that the buyers arrive is random. That is, for a random permutation σ of $[n]$, at every step $i \in [n]$, buyer $\sigma(i)$ (and therefore their distribution $D_{\sigma(i)}$) and price $X_{\sigma(i)}$ are revealed. As such, **Prophet Secretary** can be described as a generalization of **I.I.D. Prophet Inequalities** from a single distribution to n (possibly distinct) distributions.

Let f_i denote the probability density function of D_i, and F_i denote the cumulative distribution function of D_i.

Many algorithms will define thresholds $\langle \tau_1, \ldots, \tau_n \rangle$, which are the thresholds for deciding to sell the item. That is, for any buyer i, the algorithm will sell the item to i if $X_i \geq \tau_i$.

1.2 Related Work for the Prophet Secretary Problem

Prophet Secretary (2015) [3] is the original paper to suggest a marriage between the **Prophet Inequalities** problem and the **Secretary** problem[2] in optimal stopping theory. In this paper, Esfandiari et al. provide a $1 - \frac{1}{e} \approx 0.63$ approximation to the **Prophet Secretary** problem and prove that no online algorithm for **Prophet Secretary** may get better than a 0.75-approximation.[3] They also prove that a Fixed-Threshold Algorithm (FTA) — an algorithm in which the item is sold to the first buyer who pays above a single <u>fixed</u> threshold — cannot achieve better than a 0.5-approximation. *Prophet Secretary for Combinatorial Auctions and Matroids (2018)* [4] approaches our same single-choice **Prophet Secretary** problem with a continuous relaxation and so offers two elegant proofs of a $1 - \frac{1}{e}$ approximation — one with a continuous discount function and one with an FTA. They surpass the impossibility result — the 0.5-approximation upper bound on FTA algorithms — found by Esfandiari et al. in the aforementioned paper by considering the **Prophet Secretary** in the continuous setting. *Prophet Secretary: Surpassing the $1 - 1/e$ Barrier (2017)* [5] eliminates the possibility that $1 - \frac{1}{e}$ is a tight bound for **Prophet**

[1] *Posted Price Mechanisms for a Random Stream of Customers (2017)* [1] show that the optimal auction, as proven by Myerson in *Optimal Auction Design (1981)* [2], is the proper benchmark to compare an online algorithm instead of the expectation of the maximum valuation. Their example: a single customer with valuation in $[1, +\infty)$ distributed according to CDF: $F(v) = 1 - 1/v$. Most of the papers, however, do not consider this counterexample, so we will present papers as they are written.

[2] The **Secretary** problem is where you have n buyers with no known distributions. Buyers reveal themselves sequentially and you must decide whether to sell to them immediately or forever pass on selling to them.

[3] More specifically, they prove that no online algorithm can achieve a competitive ratio of $0.75 + \varepsilon$

Secretary by providing a $\left(1 - \frac{1}{e} + \frac{1}{400}\right)$-approximation. *Beating* $1 - \frac{1}{e}$ *for Ordered Prophets (2017)* [6] proves that there can be no approximation with a competitive ratio better than $\frac{11}{15}$ for deterministic, distribution-insensitive algorithms. *Prophet Secretary Through Blind Strategies (2018)* [7] improves on this impossibility bound to a factor of 0.732 and themselves find an algorithm for a 0.669-approximation.

Paper (Year)	Lower Bound	Upper Bound
Prophet Secretary (2015) [3]	$1 - 1/e$	0.75
Prophet Secretary: Surpassing the $1 - 1/e$ Barrier (2017) [5]	$1 - 1/e + 1/400$	$11/15 \approx 0.733^{\dagger}$
Prophet Secretary for Combinatorial Auctions and Matroids (2018) [4]	$1 - 1/e$	—
Prophet Secretary Through Blind Strategies (2019) [7]	0.669	$\sqrt{3} - 1 \approx 0.732$

Table 1: Summary of results for the Prophet Secretary problem.

† For deterministic, distribution-insensitive algorithms.

1.3 Related Work for the I.I.D. Prophet Inequalities Problem

Comparisons of Stop Rule and Supremum Expectations of I.I.D. Random Variables (1982) [8] supplies a $\left(1 - \frac{1}{e}\right)$-approximation for **I.I.D. Prophet Inequalities** and conjecture that the upper-bound for approximation algorithms (for arbitrarily large n) is $\frac{1}{1+\frac{1}{e}} \approx 0.731$. *Beating* $1 - \frac{1}{e}$ *for Ordered Prophets (2017)* [6] refutes the 35-year-old open conjecture by finding an algorithm for a 0.738-approximation. *Posted Price Mechanisms for a Random Stream of Customers (2017)* [1] extends an approach from a similar problem (sequential posted pricing) to find a 0.745-approximation and proves that no algorithm can do better.

Paper (Year)	Lower Bound	Upper Bound
Comparisons of Stop Rule and Supremum Expectations of I.I.D. Random Variables (1982) [8]	$1 - 1/e$	0.731 (conjecture)
Beating $1 - 1/e$ for Ordered Prophets (2017) [6]	0.738	—
Posted Price Mechanisms for a Random Stream of Customers (2017) [1]	0.745	0.745

Table 2: Summary of results for the I.I.D. Prophet Inequalities problem.

1.4 This Paper

In this paper, we review many of the above referenced papers and examine the possibility of extension from **I.I.D. Prophet Inequalities** algorithms to the **Prophet Secretary** problem. Specifically, we present an abridged review of the **Prophet Secretary** results in the papers [3, 4, 5, 7] mentioned above. In the course of doing so, we suggest why a certain type of extension of the continuous algorithm presented by Ehsani et al. [4] will not produce a better guarantee. Then, we present the **I.I.D. Prophet Inequalities** algorithm shown by Alobhasani et al. [6] and attempt to extend it to a solution for the **Prophet Secretary** problem. Subsequently, we present the optimal the **I.I.D. Prophet Inequalities** online algorithm found by Correa et al. [1] and conclude by offering summarized thoughts on the intuition of extending algorithms for the **I.I.D. Prophet Inequalities** to the **Prophet Secretary** problem.

2 Prophet Secretary (2015) [3]

In this paper, Esfandiari et al. proves both a threshold-based algorithm for a $\left(1 - \frac{1}{e}\right)$-approximation and an impossibility result — that no online single-fixed-threshold algorithm may do better than a $\left(\frac{1}{2} + \frac{1}{2n}\right)$-approximation.

2.1 Algorithm for $\left(1 - \frac{1}{e}\right)$-Approximation

This algorithm achieves a $\left(1 - \frac{1}{e}\right)$-approximation by selling to the first buyer i such that $X_{\sigma(i)} \geq \tau_i$. They construct n thresholds τ_i by:

1. $\tau_i = \alpha_i \cdot \text{OPT}$
2. $\alpha_n = \frac{1}{n+1}$
3. $\alpha_i = \frac{1+n\alpha_{i+1}}{n+1}$ for $i \in [n-1]$

Let z_i be the profit from round i ($z_i = X_{\sigma(i)}$ if item $X_{\sigma(i)}$ is chosen, and $z_i = 0$ otherwise) and $\theta(i)$ be the probability that the algorithm doesn't choose a sample from the first i samples. The two important lemmas they show are:

Lemma 1.

$$\sum_{i=1}^{n} \int_{0}^{\tau_i} \Pr[z_k \geq x]\, dx \geq OPT \cdot \sum_{i=1}^{n} (1 - \theta(i))(\alpha_i - \alpha_{i+1})$$

Lemma 2.

$$\sum_{i=1}^{n} \int_{\tau_i}^{\infty} \Pr[z_k \geq x]\, dx \geq OPT \cdot \sum_{i=1}^{n} \frac{\theta(i)}{n}(1 - \alpha_i)$$

Using these lemmas, they show that the expected profit is at least

$$OPT \cdot \left(\alpha_1 + \sum_{i=1}^{n} \theta(i) \left(\frac{1}{n} - \frac{\alpha_i}{n} - \alpha_i + \alpha_{i+1} \right) \right)$$

They set

$$\alpha_i = \sum_{k=0}^{n-k} \frac{n^k}{(1+n)^{k+1}}$$

so that the summation goes to 0, and they show that as $n \to \infty$, $\alpha_1 \to 1 - 1/e$, giving a competitive ratio of $1 - 1/e \approx 0.63$.

3 Prophet Secretary for Combinatorial Auctions and Matroids (2018) [4]

This paper by Ehsani et al. conceptualizes the buyers arriving in a random order by considering them as arriving at a time T_i uniformly at random in $[0, 1]$. They define[4] the value of an algorithm $ALG = P + R$, where P is the price at which the algorithm was willing to sell the item and R is the difference between the price paid and what the algorithm would have been willing to sell the item for. P is therefore defined as 0 if the algorithm did not sell the item, or otherwise $\tau(t)$, which is the algorithm's threshold function that denotes the threshold above which the algorithm would be willing to sell the item at any time t. R is therefore defined as 0 if the algorithm did not sell the item, or otherwise the price paid minus $\tau(t)$.

Formally, $P = \begin{cases} \tau(t), & \text{algorithm sold item} \\ 0, & \text{otherwise} \end{cases}$ and $R = \begin{cases} ALG - \tau(t), & \text{algorithm sold item} \\ 0, & \text{otherwise} \end{cases}$

3.1 FTA for Prophet Secretary

Assuming without loss of generality all distributions have finite expectation and a continuous CDF, Ehsani et al. sets a single threshold τ to be any value such that $\mathbb{P}[\max_{i=1,\dots,n} X_i \geq \tau] = 1 - \frac{1}{e}$.

They show

$$\mathbb{E}[R] \geq \mathbb{E}[(\text{OPT} - \tau)^+] \cdot \int_{t=0}^{1} \mathbb{P}[\text{item is unsold at time } t]\, \mathrm{d}t$$

[4]Their formulation of the setting is slightly different: they wish to sell the item to the buyer with the greatest value, whereas most other setups want to maximize seller's revenue. That said, the fundamentals are identical to other cases. We will use slightly different notation than they suggested to keep this paper consistent.

and bound $\mathbb{P}[\text{item is unsold at time } t] \geq e^{-t}$ (using the definition of τ) in order to get

$$\mathbb{E}[ALG] = \mathbb{E}[P] + \mathbb{E}[R] \geq \left(1 - \frac{1}{e}\right)\tau + \mathbb{E}[(\text{OPT} - \tau)^+] \cdot \left(1 - \frac{1}{e}\right) = \left(1 - \frac{1}{e}\right)\mathbb{E}[\text{OPT}]$$

They believe that the continuous setting allows them to surpass the bound set in the *Prophet Secretary (2015)* [3] paper. Note that they then prove that there can be no FTA with a guarantee better than a $\left(1 - \frac{1}{e}\right)$-approximation.

3.2 Algorithm for $1 - \frac{1}{e}$ Approximation from Continuous Analysis

The algorithm arrives at a $\left(1 - \frac{1}{e}\right)$-approximation by setting the threshold for selling the item to a buyer that arrives at any time t to be $\tau(t) := \text{OPT} \cdot (1 - e^{t-1})$.

Definitions: Let $\alpha(t) : [0, 1] \to [0, 1]$ be a differentiable discount function. Let a function $r(t) : [0, 1] \to \mathbb{R}_{\geq 0}$ be a *residual* if it satisfies the following for all α:

1. $r(0) = \text{OPT}$ (recall $\text{OPT} = \mathbb{E}[\max_{i=1,\ldots,n} X_i]$)
2. $\mathbb{E}[P] \geq -\int_{t=0}^{1} \alpha(t) \cdot r'(t)\,\mathrm{d}t = \int_{t=0}^{1} r(t) \cdot \alpha'(t)\,\mathrm{d}t - [r(t) \cdot \alpha(t)]_{t=0}^{1}$
3. $\mathbb{E}[R] \geq \int_{t=0}^{1} (1 - \alpha(t)) \cdot r(t)\,\mathrm{d}t$

They show a $\left(1 - \frac{1}{e}\right)$-approximation by proving that $r(t) := \text{OPT} \cdot \mathbb{P}[\text{item not sold before } t]$ is a residual function, setting $\alpha(t)$ appropriately, and using these two facts to have to bound $\mathbb{E}[ALG]$.

We omit the details of proving that $r(t)$ is a residual function, except to note that they do not make any assumptions about the distributions $\{D_1, \ldots, D_n\}$ in showing this. Using our bounds on R and P from the definition of a residual,

$$\begin{aligned}
\mathbb{E}[ALG] &= \mathbb{E}[R] + \mathbb{E}[P] \\
&\geq \int_{t=0}^{1} r(t) \cdot (1 - \alpha(t))\,\mathrm{d}t - [r(t)\alpha(t)]_{t=0}^{1} + \int_{t=0}^{1} r(t)\alpha'(t)\,\mathrm{d}t \\
&= \int_{t=0}^{1} r(t) \cdot (1 - \alpha(t) + \alpha'(t))\,\mathrm{d}t - [r(t)\alpha(t)]_{t=0}^{1} \qquad (*)
\end{aligned}$$

By choosing $\alpha(t)$ such that $1 - \alpha(t) - \alpha'(t) = 0$ — specifically, let $\alpha(t) = 1 - e^{t-1}$ — they can arrive at

$$\mathbb{E}[ALG] \geq [-r(1) \cdot (1 - e^0)] + [r(0) \cdot (1 - \frac{1}{e})] = 0 + \text{OPT} \cdot \left(1 - \frac{1}{e}\right)$$

3.3 New Result: Impossibility of Extension

We believe, just as the paper asserts, that reasoning about $r(t)$ is very difficult. Therefore, in exploring the possibility of extending this algorithm, we tried to find a different definition of $\alpha(t)$ that would give a better approximation (from equation $(*)$) by maximizing $(1 - \alpha(t) + \alpha'(t))$ and keeping $[r(t)\alpha(t)]_{t=0}^{1}$ constant. More specifically, that would require:

1. $\alpha(1) = 0$
2. $\alpha(0) \geq 1 - \frac{1}{e}$
3. $1 - \alpha(t) + \alpha'(t) \geq a, \quad \forall t \in [0, 1], \; a > 0$

Note that $\forall t \in [0, 1]$, we have $r(t) \geq 0$ and $0 \leq \alpha(t) \leq 1$.

Solving the differential equation of $1 - \alpha(t) + \alpha'(t) \geq a$, we have $\alpha(t) = 1 - a + ce^t$ for any $c < 0$. We can get a closed form for c from the first condition: $\alpha(1) = 1 - a + ce^1 \implies c = \frac{a-1}{e}$. This means $\alpha(0) = (1 - a) + \frac{a-1}{e} \cdot e^0 = (1 - a)\left(1 - \frac{1}{e}\right)$, which only fulfills the second condition at $a = 0$.

4 Prophet Secretary: Surpassing the $1 - 1/e$ Barrier (2017) [5]

This paper by Azar et al. aims to show that $1 - 1/e$ is not a tight lower bound for the **Prophet Secretary** problem. They achieve this by demonstrating an algorithm which achieves a $1 - 1/e + 1/400$-approximation, using similar ideas to the original 2015 paper by Esfandiari et al.
First, let $X = \max_i X_i$ and scale the distributions so $\mathbb{E}[X] = 1$. Recall that

$$\mathbb{E}[X] = \int_0^\infty \Pr[X \geq x]\, dx$$

They define the contribution of a particular interval $I \subseteq \mathbb{R}$ to $\mathbb{E}[X]$ as:

$$\int_{x \in I} \Pr[X \geq x]\, dx$$

They then improve the bound by considering three types of cases.

1. Cases where the contribution of $[0, 1 - 1/e]$ to $\mathbb{E}[X]$ is small.
2. Cases where more than one of the X_i's exceeds a certain threshold.
3. The rest of the cases.

We can formally define these cases. First, let us have four parameters, independent of n and the distributions: $a \in [0, 1-1/e], b > 1, c \in [0,1], d \in [0,1]$. These parameters are chosen so that they satisfy $ac + d > 1$, and their values will be set at the end. Furthermore, let T be defined so

$$\int_0^T \Pr[X > x]\, dx = 1 - d$$

This implies that $T \geq 1-d$, since $\Pr[X > x] \leq 1$. Finally, we define each case, as well the the appropriate thresholds τ_i, where we choose X_i if $X_i \geq \tau_i$.

1. If $\int_0^a \Pr[X \geq x]\, dx \leq c \cdot a$, then

$$\tau_i = \begin{cases} 1 - e^{(i-1)/n-1} & \text{if } \frac{i-1}{n} > 1 + \ln(1+a), \\ \frac{1}{c} - \left(\frac{1}{c} - a\right)\left(\frac{e^{(i-1)/n-1}}{1-a}\right)^c & \text{otherwise.} \end{cases}$$

2. Else if $\sum_{i=1}^n \Pr[X_i \geq T] \geq b$, then

$$\tau_i = \begin{cases} \left(1 - e^{b((i-1)/n-1)}\right) \cdot \frac{b+d-bd}{b} & \text{if } \frac{i-1}{n} > 1 - \frac{1}{b}\ln\frac{b+d-db}{b}, \\ 1 - d \cdot \left(\frac{b+d-bd}{b}\right)^{1/b} \cdot e^{(i-1)/n-1} & \text{otherwise.} \end{cases}$$

3. Else, for all i, $\tau_i = T$.

All that remains is to analyze each case. Then, let z_i be the profit from round i ($z_i = X_{\sigma(i)}$ if item $X_{\sigma(i)}$ is chosen, and $z_i = 0$ otherwise). Then, the expected profit from an algorithm with thresholds τ_i is

$$\sum_{i=1}^n \int_0^\infty \Pr[z_i \geq x]\, dx = \sum_{i=1}^n \int_0^{\tau_i} \Pr[z_i \geq x]\, dx + \sum_{i=1}^n \int_{\tau_i}^\infty \Pr[z_i \geq x]\, dx$$

They recall two lemmas from Esfandiari et al.:

Lemma 1.

$$\sum_{i=1}^n \int_0^{\tau_i} \Pr[z_i \geq x]\, dx = \sum_{i=1}^n (1 - \theta(i))(\tau_i - \tau_{i+1}) = \tau_1 - \sum_{i=1}^n \theta(i)(\tau_i - \tau_{i=1})$$

Lemma 2. $\Pr[X_i \geq x] \geq \frac{\theta(i)}{n}\Pr[X \geq x]$ for $x \geq \tau_i$, and

$$\int_{\tau_i}^\infty \Pr[z_i \geq x]\, dx \geq \frac{\theta(i)}{n}\left(1 - \int_0^{\tau_i} \Pr[X \geq x]\, dx\right) \geq \frac{\theta(i)}{n} \cdot (1 - \tau_i)$$

where $\theta(i)$ is the probability that the algorithm doesn't choose a sample from the first i samples.

4.1 Case 1

They show the following lemma for i where $\frac{i-1}{n} \leq 1 + \ln(1+a)$.

Lemma 3. If $\frac{i-1}{n} \leq 1 + \ln(1+a)$, then

$$\int_{\tau_i}^{\infty} \Pr[z_i \geq x]\, dx \geq \frac{\theta(i)}{n} \cdot (1 - c \cdot \tau_i)$$

They use Lemma 3 to bound $\int_{\tau_i}^{\infty} \Pr[z_i \geq x]\, dx$ for $\frac{i-1}{n} \leq 1 + \ln(1+a)$ and Lemma 2 to bound it for $\frac{i-1}{n} > 1 + \ln(1+a)$. This gives a stronger bound than using Lemma 2 alone. We bound $\sum_{i=1}^{n} \int_0^{\tau_i} \Pr[z_i \geq x]\, dx$ using Lemma 1. Then, they show that this gives the algorithm an expected profit of at least

$$\frac{1}{c} - \left(\frac{1}{c} - a\right)(e(1-a))^{-c} - \frac{\gamma_1}{n}$$

where γ_1 is an absolute constant.

4.2 Case 2

They show the following lemma for i where $\frac{i-1}{n} \geq 1 - \frac{1}{b} \ln \frac{b+d-db}{b}$.

Lemma 4. If $\frac{i-1}{n} \geq 1 - \frac{1}{b} \ln \frac{b+d-db}{b}$, then

$$\int_{\tau_i}^{\infty} \Pr[z_i \geq x]\, dx \geq \frac{\theta(i)}{n} \cdot (b + d - bd - b \cdot \tau_i)$$

They use Lemma 4 to bound $\int_{\tau_i}^{\infty} \Pr[z_i \geq x]\, dx$ for k where $\frac{i-1}{n} \geq 1 - \frac{1}{b} \ln \frac{b+d-db}{b}$, and Lemma 2 to bound it for k where $\frac{i-1}{n} < 1 - \frac{1}{b} \ln \frac{b+d-db}{b}$. Again, this gives a better bound than Lemma 2 alone. They bound $\int_0^{\tau_i} \Pr[z_i \geq x]\, dx$ using Lemma 1, similar to before. They then show that this gives an expected profit of at least

$$1 - \frac{d}{e} \cdot \left(\frac{b + d - bd}{d}\right)^{1/b} - \frac{\gamma_2}{n}$$

where γ_2 is an absolute constant.

4.3 Case 3

We will go through the general idea behind this proof, but we will skip the math that goes into it because it is quite involved. Recall that in this case, we just set all thresholds to a value T.

1. They show that there exists some X_i (WLOG, let this be X_1) that is larger than the rest of the variables with high probability, and that it has a high enough expectation.
2. They prove that the algorithm is very unlikely to pick a sample before reaching X_1, so it gets most of the expected value from X_1.
3. The algorithm sees $1/2$ of the other X_i's before seeing X_1, on average. Even if one of the X_i's beats the threshold, we still get our value about half the time.

They use a similar method as above, bounding each term separately, to show each of these parts. They find that this gives an expected profit of at least

$$h(1-d) + \left(\frac{1}{2} - \frac{b-g}{3}\right) \cdot (ca - (1-d) - (a - 1 + d)(b - g)) + \frac{1}{2}\left(1 - \frac{b-g}{3}\right) \cdot d$$

where $h = \frac{ca-1+d}{a-1+d}$ and $g \in (0, 1]$ is such that $(1-g)^{1/g} = (1-h)^{1/b}$.

4.4 Overall Competitive Ratio

They find that setting $a = 1 - 1.31/e$, $c = 0.98$, $d = 0.62$, and $g = 0.88$ gives the desired competitive ratio. We can check that $ac + d > 1$, and we get $h \approx 0.925$ and $b \approx 1.075$. Then, the lower bounds from all three cases are at least $c^* - O(n^{-1})$ for some $c^* > 1 - 1/e + 1/400$. They then get rid of the $O(n^{-1})$ term by showing that they can find an algorithm that achieves a $c^* - \varepsilon$-approximation for $c^* > 1 - 1/e + 1/400$.

4.5 Extensions

Because this paper was so focused on just beating the $1 - 1/e$ bound, the algorithm they come up with is a little contrived and difficult to build on or extend. As a result, we didn't end up using this paper for much; it was only significant as the first result that beat the $1 - 1/e$ bound.

5 Prophet Secretary Through Blind Strategies (2019) [7]

This paper by Correa et al. shows a 0.669-approximation for the **Prophet Secretary** problem, using a multi-threshold algorithm which they call a "blind strategy." Intuitively, a blind strategy is a strategy that designs thresholds independently of what distributions the bidders have, and is conditioned only on the expected maximum value. They define a blind strategy as the following:

1. Let $\alpha : [0, 1] \to [0, 1]$ be a non-increasing function.
2. Draw $q_1, ..., q_n$ from a uniform distribution on $[0, 1]$. Without loss of generality, we assume $q_1 \leq q_2 \leq ... \leq q_n$. Define thresholds τ_i as

$$\Pr[\max_{i\in[n]} X_i \leq \tau_i] = \alpha(q_i)$$

3. Then, we go through the bidders, who are ordered in some permutation σ, and we stop and pick $X_{\sigma(i)}$ when $X_{\sigma(i)} > \tau_i$.

They let $\alpha_i = \alpha(i/n)$, and they prove the following lemma.

Lemma 1. Fix α, which fixes $\alpha_1, ..., \alpha_n$. Consider the sequence of thresholds such that

$$\Pr[\max_{i\in[n]} X_i \leq \tau_i] = \alpha_i$$

Let T be the stopping time of the algorithm for these thresholds. Then, for all $k \in [n]$, we have

$$\frac{1}{n}\sum_{j\in[k]} 1 - \alpha_j \leq \Pr[T \leq k] \leq 1 - \left(\prod_{j=1}^{k} \alpha_j\right)^{1/n}$$

Using this lemma, they prove the following theorem.

Theorem 2. Let $\alpha : [0, 1] \to [0, 1]$ be a non-increasing function, and let T be the stopping time of the algorithm. Then, we have for $t > 0$,

$$\Pr[X_{\sigma(T)} > t] \geq \min_{j\in[n+1]}\{f_j(\alpha)\}\mathbb{P}\left[\max_{i\in[n]}\{X_i\} > t\right]$$

where, for all $j \in [n + 1]$ and taking $\alpha(\frac{n+1}{n}) = 0$, we have

$$f_j(\alpha) = \sum_{k=1}^{j-1} \frac{1 - \alpha(k/n)}{n(1 - \alpha(j/n))} + \frac{1}{n}\sum_{k=j}^{n}\left(\prod_{l=1}^{k}\alpha(l/n)\right)^{1/n}$$

By doing a Riemann sum analysis when $n \to \infty$, they get that

$$\lim_{n\to\infty}\min_{j\in[n+1]}\{f_j(\alpha)\} = \min\left\{\int_0^1 1 - g(y)\,dy, \inf_{x\in[0,1]}\int_0^x \frac{1 - g(y)}{1 - g(x)}\,dy + \int_x^1 \exp\left(\int_0^y \ln g(w)\,dw\right)\,dy\right\},$$

where the left-hand side is exactly the competitive ratio for this algorithm. They optimize this quantity by solving a differential equation, getting a bound of 0.665. Using the same general approach but with sharper inequalities, they show that they can improve this bound to 0.669.

They then show that no blind strategy can achieve a competitive ratio better than 0.675, and that no strategy can achieve a competitive ratio better than $\sqrt{3} - 1 \approx 0.732$. This is significant because it is the first time that a separation has been shown between the **Prophet Secretary** and the **I.I.D. Prophet Inequalities** problems.

6 Beating $1 - 1/e$ for Ordered Prophets (2017) [6]

We first describe the method used by Abolhasani et al. to obtain a 0.738-approximation for the **I.I.D. Prophet Inequalities**, then we describe how this method can be extended to the **Prophet Secretary** case. Their algorithm is as follows:

1. Let $a = 1.306$. (This is the root of $\cos(a) + \sin(a)/a - 1 = 0$.)
2. Let $\tau_i = F^{-1}(\cos(ai/n)/\cos(a(i-1)/n))$.
3. Pick the first item where $X_i \geq \tau_i$.

6.1 Algorithmic Analysis

First, they define a threshold function. A continuous and strictly decreasing function $h : [0, 1] \to [0, 1]$ where $h(0) = 1$ is a threshold function if:

1. h is strictly concave.
2. For every $\varepsilon > 0$ there exists a $\delta_0 < \varepsilon$ such that for every $\delta < \delta_0$ and $\varepsilon + \delta \leq s \leq 1$,

$$\frac{h'(s-\delta)}{h(s-\delta)} \leq (1-\varepsilon)\frac{h'(s)}{h(s)}$$

For a threshold function h, they show that the thresholds τ_i, generated by

$$\tau_i = F^{-1}\left(\frac{h(i/n)}{h((i-1)/n)}\right),$$

form a decreasing sequence. They go on to define an α-strong threshold function. A threshold function is α-strong if:

1. $h(1) \leq 1 - \alpha$
2. $\int_0^1 h(r)\, dr \geq \alpha$
3. $\forall\, 0 \leq s \leq 1, 1 - h(s) - \frac{h'(s)}{h(s)} \int_0^1 h(r)dr \geq \alpha \left(1 - \exp\left(\frac{h'(s)}{h(s)}\right)\right)$

They show that if a threshold function h is α-strong, then using the above algorithm will give an α-approximation for the **I.I.D. Prophet Secretary**. Finally, they demonstrate that $h(i/n) = \cos(ai/n)$ is an α-strong threshold function for $\alpha = 1 - cos(a) = 0.7388$. Therefore, their algorithm achieves a 0.7388-approximation for the **I.I.D. Prophet Inequalities** problem.

6.2 New Result: Extension to Prophet Secretary

Now, we wish to extend this idea to the **Prophet Secretary** problem. We take inspiration from Correa et al. (2019) [7], who show a 0.669-approximation for the **Prophet Secretary** problem using a similar decreasing threshold approach. Recall that they design thresholds τ_i using a function $g(i/n)$ such that

$$\mathbb{P}\left[\max_{i=1,\ldots,n} X_i \leq \tau_i\right] = g(i/n)$$

The quantity on the LHS is essentially the **Prophet Secretary** analog of $F(\tau_i)$ in the earlier expression (rearranged):

$$F(\tau_i) = \mathbb{P}[X_i \leq \tau_i] = \frac{h(i/n)}{h((i-1)/n)}$$

Therefore, it seems reasonable to let

$$g(i/n) = \frac{h(i/n)}{h((i-1)/n)} = \frac{\cos(ai/n)}{\cos(a(i-1)/n)}$$

The approximation guarantee shown by Correa et al. [7] for the function g is

$$\min\left\{\int_0^1 1-g(y)\,dy,\ \inf_{x\in[0,1]}\int_0^x \frac{1-g(y)}{1-g(x)}\,dy+\int_x^1 \exp\left(\int_0^y \ln g(w)\,dw\right)\,dy\right\},$$

so all that's left to do is to compute this value. Letting $n=1000$, we can numerically evaluate:

$$\int_0^1 1-g(y)\,dy = 0.00133855$$

$$\inf_{x\in[0,1]}\int_0^x \frac{1-g(y)}{1-g(x)}\,dy+\int_x^1 \exp\left(\int_0^y \ln g(w)\,dw\right)\,dy = 0.278597$$

The first value seems abnormally small, so we investigate possible differences between the function g we designed and the function used by Correa et al. Although they do not provide a concrete function for their best bound, they find that using $a(y)=0.53-0.38y$ gives a guarantee of $0.657>1-1/e$. We plot both functions in Figure 1.

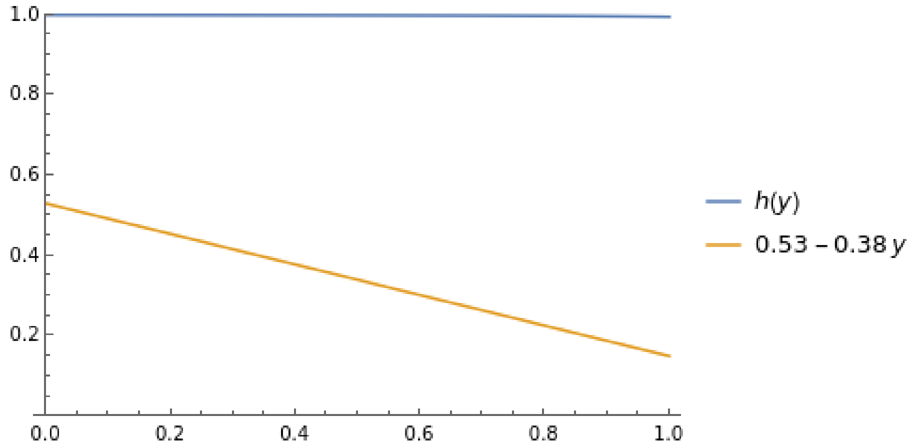

Figure 1: Plots of both functions.

We quickly notice that the thresholds from the i.i.d. problem are extremely high compared to the thresholds for the **Prophet Secretary**. In fact, we realize that as $n\to\infty$, $g(y)\to 1$ for all $y\in[0,1]$. As $n\to\infty$ in the i.i.d. problem, the probability that a very large draw occurs increases as well. Therefore, for very large n, an algorithm for the i.i.d. problem can be extremely greedy and wait for a huge value to stop on. However, for the **Prophet Secretary** problem, the distributions of each bidder are different. The likelihood that you eventually get a draw that is far above the expected value of a bidder's distribution similarly increases as $n\to\infty$. However, we are not guaranteed that the expected value of this distribution is large, so this draw may not be greater than a very high threshold. As a result, we should be more conservative with our thresholds, and take values that aren't nearly as good.

This gives one possible idea: rescale the above threshold by $1/e$ to make it more conservative. This gives

$$g(y)=\frac{1/e\cdot\cos(ay)}{\cos(a(y-1/n))}$$

We can compute the bound numerically again for $n=1000$,

$$\int_0^1 1-g(y)\,dy = 0.632613$$

$$\inf_{x\in[0,1]}\int_0^x \frac{1-g(y)}{1-g(x)}\,dy+\int_x^1 \exp\left(\int_0^y \ln g(w)\,dw\right)\,dy = 0.631951$$

So this function also does not end up beating $1-1/e$, but gets closer. However, this function is close to constant, and doesn't decrease the thresholds as we get farther in, which may be why it does not perform as well.

7 Posted Price Mechanisms for a Random Stream of Customers (2017) [1]

This paper by Correa et al. extends a result from posted price mechanisms to **I.I.D. Prophet Inequalities**. Posted price mechanisms set prices for each buyer, unlike in **Prophet Inequalities** where the buyer suggests the price. Correa et al. discuss multiple types of such mechanisms, but the one most similar to **I.I.D. Prophet Inequalities** is the *adaptive* posted price mechanism. In this mechanism, the buyers arrive in random order and immediately reveal their distribution. The algorithm then chooses a price for the object and the buyer can either buy the item at that price or pass on them forever.

Correa et al. explicitly extends their results from the adaptive posted price mechanism context to the **I.I.D. Prophet Inequalities** problem to achieve a 0.745-approximation.

7.1 0.745-Approximation

First, define constants β and y:

$\beta \approx 0.745$ is defined by $1 = \int_0^1 \left(y(1-\ln(y)) + \beta - 1\right)^{-1} \, \mathrm{d}y$.

y is the unique solution to the ordinary differential equation: $y' = y \cdot \ln(y) - y - \frac{1}{\beta} + 1$ and $y(0) = 1$.

Then, choose, independent of the buyers' distributions, $0 = \varepsilon_0 < \varepsilon_1 < \cdots < \varepsilon_{n-1} < \varepsilon_n = 1$ such that $\rho_1 = \cdots = \rho_n$ with ρ_i defined by $\rho_1 = \frac{1}{\alpha_1}$ and $\rho_{i+1} = \frac{\rho_i}{\alpha_{i+1}} \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-1} \, \mathrm{d}q$ for $i = 1, \ldots, n-1$ and $\alpha_i = \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-r)^{n-2} \, \mathrm{d}r$. As $n \to \infty$, it has been proven that $\varepsilon_i \to 1 - y\left(\frac{i}{n}\right)^{1/(n-1)}$.

Following the de-randomization scheme of this algorithm proposed in *Prophet Inequalities for I.I.D. Random Variables from an Unknown Distribution (2021)* [9], we can let

$$q_i = \int_{\varepsilon_{i-1}}^{\varepsilon_i} \frac{(n-1)q(1-q)^{n-2}}{\alpha_i} \mathrm{d}q$$

Finally, we let the threshold $\tau_i = F^{-1}(1 - q_i)$.

This algorithm gets expected revenue equal to $\dfrac{1}{n\alpha_1}\mathbb{E}\left[\max_{i=1,\ldots,n} X_i\right] \geq \beta^{-1}\mathbb{E}\left[\max_{i=1,\ldots,n} X_i\right]$.

Importantly, this paper proves that this approximation-bound is tight.

8 Conclusion

8.1 Intuitions on Extensions from I.I.D. Prophet Inequalities to Prophet Secretary

The first thing to note about any extension from **I.I.D. Prophet Inequalities** to **Prophet Secretary** is that the upper and lower bounds for the two problems are different: Correa et al. [1] found a 0.669-approximation, which is the best currently known online algorithm for **Prophet Secretary** and proved that no algorithm can achieve better than a 0.732-approximation; in contrast, a 0.745-approximation has been found for **I.I.D. Prophet Inequalities** and shown to be tight. Therefore, we expect most online algorithms for **I.I.D. Prophet Inequalities** that take advantage of the uniform distributions will not be as successful in their **Prophet Secretary** extensions.

Intuitively, the benefit of having I.I.D. distributions is that there is no situation where the algorithm would need to wait for a distribution that is much larger than the other distributions in expectation. Instead, it only needs to wait until it gets a draw from the same distribution which is far above average. If there are n (possibly distinct) distributions, a **Prophet Secretary** algorithm might need to wait for the distribution with a large expected value, and it might need to get an above average draw from this distribution too. For example, if $D_1 = \begin{cases} 1000, & \text{w.p. } \frac{1}{2} \\ 1, & \text{w.p. } \frac{1}{2} \end{cases}$ and $D_2 = D_3 = \cdots = D_n = \begin{cases} 10, & \text{w.p. } \frac{1}{2} \\ 100, & \text{w.p. } \frac{1}{2} \end{cases}$, then the optimizing algorithm will wait for until buyer $\sigma(t) = 1$ arrives so that it can increase its expected value to at least 500.5. Any threshold function that does not consider the still unseen buyers' distributions will therefore not achieve an optimal online algorithm for **Prophet Secretary**.

We see that the approaches to the two problems often involve some series of descending thresholds, where we choose the element at time i if the value drawn, X_i, is greater than the threshold τ_i. In the **I.I.D. Prophet Inequalities** algorithms, the thresholds can get greedier and greedier as $n \to \infty$, because within these n draws, we are very likely to get a draw that is very far above the expectation.

However, in the **Prophet Secretary** case, we have no such guarantee when $n \to \infty$. Of course, we will still be very likely to find some draw with a large (CDF) value from *some* distribution, but in this case, we have no guarantee that its actual value will be large. Therefore, the thresholds for the **Prophet Secretary** problem cannot be as greedy, and specifically, they cannot tend toward 1 as $n \to \infty$.

The algorithm proposed by Correa et al. [1] that achieves a 0.745-approximation for the **I.I.D. Prophet Inequalities** designs thresholds that are insensitive both to the number of bidders and to the distributions themselves. In another paper, Correa et al. [7] prove a 0.675 upper bound on these types of algorithms that design the same thresholds regardless of what the distributions are. We believe that the key to bridging the remaining gap in the upper and lower bound lies in taking into account the knowledge we have about the distributions and which of them are still unseen. As we demonstrated above, without utilizing this knowledge, the algorithm may not be able to make the optimal decisions. However, the analysis for distribution-dependent approaches is much more complex than the existing distribution-independent approaches.

8.2 Further Extensions of the Prophet Secretary Problem

Another avenue to consider for the **Prophet Secretary** is an instance where the distributions are unknown, but we are allowed to draw samples to gain knowledge about them. Correa et al. [9] have shown that for the **I.I.D. Prophet Inequalities** problem, using $O(n^2)$ samples gives full distributional knowledge to reach the 0.745 upper bound (discounting by an additive ε). However, it is less clear how sampling would be done for the **Prophet Secretary** problem. We imagine two possible approaches:

1. We know distributions there are n (possibly distinct) distributions, and for each sample drawn, we are told from which distribution it came.

2. We know distributions there are n (possibly distinct) distributions, but for each sample drawn, we do not know the distribution from which it came.

Correa et al. [10] have shown a 0.635-approximation for the first approach, using a single sample from each distribution, but we imagine that more can be done to reach a better bound using more samples. The second approach appears far more difficult, as the distributions will be far more difficult to separate from each other.

References

[1] J. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld, “Posted price mechanisms for a random stream of customers,” in *Proceedings of the 2017 ACM Conference on Economics and Computation*, ser. EC ’17. New York, NY, USA: Association for Computing Machinery, 2017, p. 169–186. [Online]. Available: https://doi.org/10.1145/3033274.3085137

[2] R. B. Myerson, “Optimal auction design,” *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58–73, 1981. [Online]. Available: http://www.jstor.org/stable/3689266

[3] H. Esfandiari, M. Hajiaghayi, V. Liaghat, and M. Monemizadeh, “Prophet secretary,” *CoRR*, vol. abs/1507.01155, 2015. [Online]. Available: http://arxiv.org/abs/1507.01155

[4] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla, “Prophet secretary for combinatorial auctions and matroids,” *CoRR*, vol. abs/1710.11213, 2017. [Online]. Available: http://arxiv.org/abs/1710.11213

[5] Y. Azar, A. Chiplunkar, and H. Kaplan, “Prophet secretary: Surpassing the 1-1/e barrier,” *CoRR*, vol. abs/1711.01834, 2017. [Online]. Available: http://arxiv.org/abs/1711.01834

[6] M. Abolhassani, S. Ehsani, H. Esfandiari, M. Hajiaghayi, R. Kleinberg, and B. Lucier, “Beating 1-1/e for ordered prophets,” *CoRR*, vol. abs/1704.05836, 2017. [Online]. Available: http://arxiv.org/abs/1704.05836

[7] J. R. Correa, R. Saona, and B. Ziliotto, “Prophet secretary through blind strategies,” *CoRR*, vol. abs/1807.07483, 2018. [Online]. Available: http://arxiv.org/abs/1807.07483

[8] R. P. Kertz, “Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality,” *Journal of Multivariate Analysis*, vol. 19, no. 1, pp. 88–112, 1986. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0047259X86900953

[9] J. Correa, P. Dütting, F. Fischer, and K. Schewior, “Prophet inequalities for i.i.d. random variables from an unknown distribution,” in *Proceedings of the 2019 ACM Conference on Economics and Computation*, ser. EC ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 3–17. [Online]. Available: https://doi.org/10.1145/3328526.3329627

[10] J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto, “The two-sided game of googol and sample-based prophet inequalities,” 2019. [Online]. Available: http://arxiv.org/abs/1907.06001