Prophet Inequalities for I.I.D. Random Variables with Linear Samples

Kevin Feng and Catherine Yu

COS521 Final Project, Fall 2021

Abstract

We examine guarantees for a specific variant of the prophet inequality problem: given a sequence of random variables $X_1, \dots, X_n$ drawn i.i.d. from an unknown distribution $\mathcal{D}$, along with γn other i.i.d. samples from $\mathcal{D}$, find a stopping time τ with guarantee α such that for all distributions, $\mathbb{E}[X_\tau] \geq \alpha \mathbb{E}[\max\{X_1, \dots, X_n\}]$. Currently, there are tight bounds known for $0 \leq \gamma \leq 1/(e-1)$ as well as $\Omega(n)$ samples, but there remains a gap for $\gamma > 1/(e-1)$.

Contents

1 Introduction

The theory of optimal stopping is concerned with choosing a time to take an action in order to maximize reward or minimize cost given imperfect information about the future. Two popular problems in this field are the secretary problem and the prophet problem.

In the secretary problem, we are shown n distinct, non-negative numbers from an unknown range in sequential order, and the goal is to stop at one of the numbers to maximize the probability that we select the maximum of the range. This problem has a simple solution: first, discard $1/e$ fraction of the numbers; then select the first number you see that is greater than all of the discarded numbers. This guarantees that you will select the maximum with probability $1/e$, and this is the best possible solution for correlated random variables.

In the classic single-choice prophet inequality problem, we are shown n non-negative numbers in sequential order $X_1, \ldots, X_n$, and each numberX_i is an independent draw from a known distribution $\mathcal{D}_i$. Our goal is to devise a stopping rule, or algorithm, that maximizes the expected value of the number we stop on in proportion to the expected maximum value of the entire sequence. The performance of stopping rules is usually measured by their competitive ratio in comparison to a prophet, who knows all the numbers in advance and gains expected reward $\mathbb{E}[\max_i X_i]$). This problem has many variations, and one of the main variations concerns whether the distributions are distinct or identical. When the distributions are distinct, there is a tight bound of $1/2$ (cite??). When the distributions are identical, Hill and Kertz first determined a lower bound of $1-1/e \approx 0.632$ [HK82], which was eventually improved to 0.745 by [CFH$^+$17], and this is known to be tight due to an impossibility result of [HK82] and [Ker86] which implies a matching upper bound.

However, there was little known until recently about another variation of the prophet problem for both identical and non-identical distributions, which is when we assume that the distributions from which values are drawn are unknown. Instead, we are given incomplete information about the distributions in the form of offline samples from the distributions $\mathcal{D}_i$, and we can use these samples to help us determine our stopping rule.

The last variation of the prophet problem we will mention here is the presence of an adversary versus randomness in the problem. An adversary can fix the prophet inequality in different ways, such as by picking the numbers X_i in the range, by picking the distributions $\mathcal{D}_i$ which the numbers are drawn from, by deciding the order that the numbers are presented, and by deciding the reward Y_i for each item that is gained when stopping on item X_i. However, the only adversary we will generally consider is the adversary that picks the distributions $\mathcal{D}_i$ ahead of time, and all other forms of fixing will be left up to randomness, i.e. the actual numbers, the order that the numbers are revealed, and their respective rewards will all be randomly decided.

When considering the prophet problem with sampling and unknown distinct distributions, the tight bound of $1/2$ from the known distributions case can be guaranteed with a single sample from each distribution $\mathcal{D}_i$ with the Single Sample Algorithm in [RWW19]. Rubinstein et al. described the following algorithm: if $\tilde{X}_1, \ldots, \tilde{X}_n$ are n independent samples from $\mathcal{D}_1, \ldots, \mathcal{D}_n$, then simply set $\max_i\{\tilde{X}_i\}$ as the threshold for stopping, in other words, we stop when we see a number in the range that is larger than all of the samples $\tilde{X}_i$. This algorithm achieves the optimal competitive ratio of $1/2$.

In this paper, we will focus on the case of an unknown identical distribution and the work that has been done to determine bounds on the optimal competitive ratio when we are given kn samples from the distribution. In [CDFS21], the authors prove a series of lower bounds for the optimal competitive ratio, starting with a tight lower bound of $1/e$ with 0 samples and even with $o(n)$ samples, and improving this to a lower bound of $1 - 1/e$ with $n - 1$ samples. They also provide a parametric lower and upper bound for γn samples for $\gamma \geq 0$, which is equal to $\ln(2) \approx 0.693$ for rules that use at most n samples and is thus nearly tight. [KNR19] continues this work by matching the $1 - 1/e$ lower bound for $\leq n$ samples, and this lower bound has been continuously improved since the initial paper by [CDFS21] in [CCES19] and [CDF$^+$20]. Finally, [RWW19] provides an algorithm with $O(n)$ samples that achieves an optimal competitive ratio of 0.745, which is on par with the known upper bound of 0.745 and is thus a tight bound for the problem using $O(n)$ samples.

In addition to this research on the best competitive ratio that can be guaranteed to solve an i.i.d. prophet inequality with kn samples, these results can be extended to other domains, notably streaming algorithms, and they can be extended to setups that are more generalized than those presented in the aforementioned papers. [CDF$^+$20] and [CCES21] respectively expand on these applications and give insight as to other applications of the field of prophet inequalities and open questions in the field.

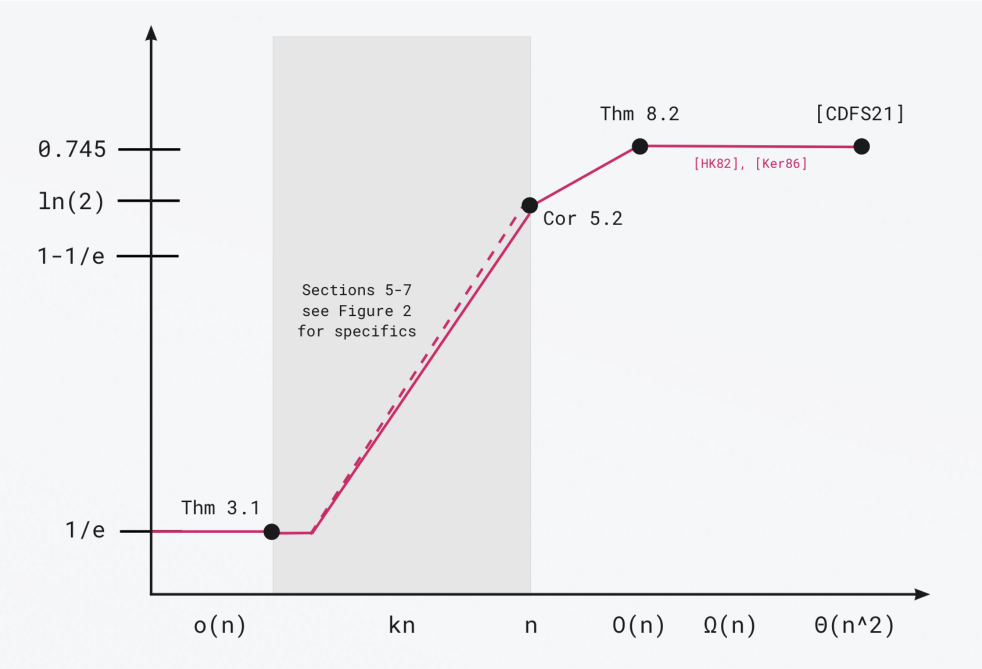

Figure 1: Overview of bounds on optimal competitive ratio for the i.i.d. prophet inequality problem with sampling labeled with the corresponding theorem, sections, or paper. The horizontal axis is the number of samples and the vertical axis is the performance guarantee. The gray rectangle represents the range for kn samples with $k \in (0, 1]$, and we discuss performance guarantees for kn samples in Sections 5-7. See Figure 2 for a zoomed in graph that shows lower bounds (solid line) and upper bounds (dashed line) more clearly based on specific papers. The results for $o(n)$ and $O(n)$ samples are tight. *Reference figure: Figure 1 in [CDFS21].*

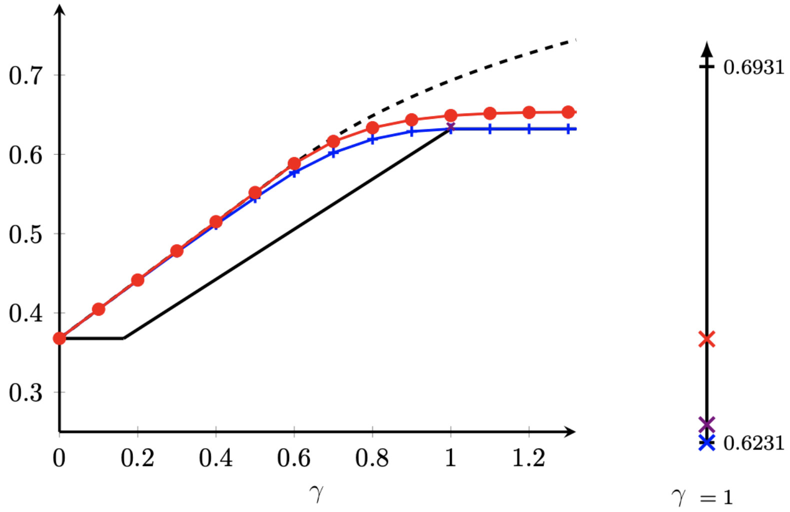

Figure 2: Visualization of the lower bounds and upper bounds established for varying γ for γn samples in the i.i.d. prophet inequality problem with sampling. The solid, red line shows the lower bound determined in [CDF$^+$20]; dashed black line and solid black line show the parametric upper bound and lower bound respectively of [CDFS21]; and the blue line denotes the lower bounds of [KNR19]. The violet mark on the right hand-side line for $\gamma = 1$ is the improved lower bound described in [CCES19]. *Reference figure: Figure 1 in [CDF$^+$20].*

2 Preliminaries

We define notation similar to that of [CDFS21] and [CDF$^+$20]. Denote $\mathbb{N}$ as the set of positive integers and $\mathbb{N}_0$ as the set of non-negative integers. For any $i \in \mathbb{N}$, we let $[i]$ be the set $\{1, \dots, i\}$.

Definition 2.1 ((k, n) stopping rule)**.** Let $k \in \mathbb{N}_0$. We consider (k, n)-stopping rules that observe k samples $S_1, \dots, S_k$, followed by n sequential random variables $X_1, \dots, X_n$, such that we decide whether or not to stop on X_i given only the samples $S_1, \dots, S_k$ and values $X_1, \dots, X_i$, for all $i \in [n]$.

Since we are concerned only with the I.I.D. case, we assume that $S_1, \dots, S_k$ and $X_1, \dots X_n$ are all i.i.d. from some distribution $\mathcal{D}$. Let f and F be its probability density function and cumulative density function, respectively. For simplicity, we assume that $\mathcal{D}$ has non-negative support and that F is continuous.

Then formally, a (k, n)-stopping rule is a family of functions $r_1, \dots, r_n$ where $r_i : \mathbb{R}_+^{k+i} \to [0, 1]$ for all $i = 1, \dots, n$ such that $r_i(s_1, \dots, s_k, x_1, \dots, x_i)$ is the probability of stopping at X_i after receiving samples $s_1, \dots, s_k$ and values $x_1, \dots, x_i$, conditioned on not stopping at $X_1, \dots, X_{i-1}$.

Definition 2.2 (Stopping time τ)**.** The stopping time τ is a random variable with support $[n] \cup \{\infty\}$ such that

$$\Pr[\tau = i \mid S_1 = s_1, \dots, X_i = x_i] = \left(\prod_{j=1}^{i-1} (1 - r_j(s_1, \dots, x_j)) \right) \cdot r_i(s_1, \dots, x_i)$$

We use the convention that $X_\infty = 0$. Then we are concerned with the expected stopping value $\mathbb{E}[X_\tau]$ and how it compares to the expected maximum $\mathbb{E}[\max\{X_1, \ldots, X_n\}]$. We say that a stopping rule achieves a competitive ratio of α if for any distribution $\mathcal{D}$, its stopping time τ satisfies $\mathbb{E}[X_\tau] \geq \alpha \cdot \mathbb{E}[\max\{X_1, \ldots, X_n\}]$.

3 A $1/e$ lower bound with 0 samples

The $1/e$ lower bound for the secretary problem immediately translates into a $1/e$ lower bound for the prophet problem with 0 samples. This is because we can simply run the same algorithm:

Algorithm 1: Secretary Algorithm

```
Data: Sequence of i.i.d. random variables $X_1, \ldots, X_n$ sampled from an unknown distribution $\mathcal{D}$
Result: Stopping time $\tau$
$T \leftarrow \max\{X_1, \ldots, X_{n/e}\}$
for $i = n/e + 1, \ldots, n$ do
    if $X_i > T$ then
        return $i$
    end
end
```

Theorem 3.1. *There exists a $(0, n)$-stopping rule that achieves a competitive ratio of $1/e$.*

Proof. By the guarantee for the secretary problem, we have that

$$\Pr[X_\tau = \max\{X_1, \ldots, X_n\}] \geq \frac{1}{e}$$

and thus we immediately have that

$$\mathbb{E}[X_\tau] \geq \Pr[X_\tau = \max\{X_1, \ldots, X_n\}] \cdot \max\{X_1, \ldots, X_n\} = \frac{1}{e} \max\{X_1, \ldots, X_n\}$$

as desired. $\square$

We note that it can be shown that $1/e$ is also an upper bound on $(0, n)$-stopping rules, and this can be extended to $(o(n), n)$-stopping rules as well. We omit the proof, but one can be found in [CDFS21].

4 A $1 - 1/e$ lower bound with $n - 1$ samples in [CDFS21]

We now show that we can do much better when given access to a linear amount of samples. The intuition is that we can use the samples, rather than a prefix of the values, to generate an appropriate threshold to choose when to stop.

In particular, first suppose we are given access to $n(n-1)$ i.i.d. samples from $\mathcal{D}$, split into n subsets $T_1, \ldots, T_n$ of $n-1$ samples. Then for each sequential X_i, we can use $\max T_i$ as a threshold value, such that we accept X_i if $X_i > \max T_i$.

But if we accept this value, then X_i is the maximum of X_i and $n-1$ other i.i.d. samples, so we have $\mathbb{E}[X_i] = \mathbb{E}[\max\{X_1, \ldots, X_n\}]$. Then the probability of termination on any X_i is independently $1/n$, so the total probability of termination is $1-(1-1/n)^n \geq 1-1/e$.

However, we have to be a bit more clever when we only have $n-1$ samples. The idea is that instead of requiring $n-1$ new samples for each X_i, we can instead choose a random subset T_i of size $n-1$ of the $n-1$ samples and $i-1$ values we have seen so far, and then stop at X_i if and only if $X_i > \max T_i$. The key lemma is as follows.

Lemma 4.1. *Conditioned on reaching X_i, the distribution of the set $\{S_1, \ldots, S_{n-1}, X_1, \ldots, X_{i-1}\}$ is identical to $n+i-2$ i.i.d. samples from $\mathcal{D}$.*

Proof. We induct on i. This holds for $i=1$ by assumption.

Now suppose this holds for $i=k-1$. That is, upon reaching X_{k-1}, the set $S_{k-1} = \{S_1, \ldots, X_{k-2}\}$ is identical to $n+k-3$ i.i.d. samples. Let T_{k-1} be the random subset of $n-1$ samples chosen from S_{k-1}. Then in order to reach X_k, we must discard X_{k-1}, i.e. we must have $X_{k-1} \leq \max T_{k-1}$. But since X_{k-1} is also i.i.d., this is just $1-1/n$. In particular, note that this probability is independent of the value of X_{k-1}.

Finally consider $S_k = \{S_1, \ldots, X_{k-1}\}$. By itself, this is clearly identical to $n+k-2$ i.i.d. samples, and we have just shown that reaching X_k is independent of the value of X_{k-1}. Thus, even when conditioned on reaching X_k, S_k is still identical to $n+k-2$ i.i.d. samples, as desired. ☐

We now present the algorithm.

Algorithm 2: Threshold Generation with $n-1$ Samples

```
Data: Sequence of i.i.d. random variables $X_1, \ldots, X_n$ sampled from an unknown distribution
      $\mathcal{D}$, with sample access to $\mathcal{D}$
Result: Stopping time $\tau$
$S_1, \ldots, S_{n-1} \leftarrow n-1$ samples from $\mathcal{D}$
$S \leftarrow \{S_1, \ldots, S_{n-1}\}$
for $i = 1, \ldots, n$ do
    if $X_i > \max S$ then
        return $i$
    else
        $S \leftarrow$ random subset of size $n-1$ from $\{S_1, \ldots, S_{n-1}, X_1, \ldots, X_i\}$
    end
end
```

Theorem 4.2. *There exists an $(n-1, n)$-stopping rule that achieves a competitive ratio of $1-(1-1/n)^n$.*

Proof. Let T_i be the random subset S at step i. We have that

$$\mathbb{E}[X_\tau] = \sum_{i=1}^{n} \mathbb{E}[X_\tau | \tau = i] \cdot \Pr[\tau = i]$$

First, by the lemma, X_τ is the maximum of n i.i.d. values from $\mathcal{D}$ independent of the value of τ, so we have that $\mathbb{E}[X_\tau | \tau = i] = \mathbb{E}[\max\{X_1, \ldots, X_n\}]$. Plugging this in, we get that

$$\mathbb{E}[X_\tau] = \sum_{i=1}^{n} \mathbb{E}[\max\{X_1, \ldots, X_n\}] \cdot \Pr[\tau = i] = (1 - \Pr[\tau \notin [n]]) \cdot \mathbb{E}[\max\{X_1, \ldots, X_n\}]$$

Again by the lemma, we have that

$$\Pr[\tau \notin [n]] = \left(1 - \frac{1}{n}\right)^n$$

which gives that

$$\mathbb{E}[X_\tau] = \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \cdot \mathbb{E}[\max\{X_1, \ldots, X_n\}]$$

as desired. □

5 A parametric lower bound with γn samples for $\gamma \leq 1$ in [CDFS21]

The result from the previous section easily translates to a parametric lower bound when we have γn samples for $\gamma \in [0, 1]$. The idea is to simply interpret the first few values in $X_1, \ldots, X_n$ as samples such that there are an equal number of samples and values again, and then apply the algorithm from the previous section.

Corollary 5.1. *There exists a* $(\gamma n, n)$*-stopping rule that achieves a competitive ratio of* $(1+\gamma)/2 \cdot (1 - 1/e)$.

Proof. Without loss of generality suppose that $\gamma n + n$ is even. Let $n' = \dfrac{1+\gamma}{2} n$, and redefine samples $S_1, \ldots, S_{\gamma n}$ and values $X_1, \ldots, X_n$ into samples

$$S'_i = S_i, \forall i \in [\gamma n];\ S'_{\gamma n + j} = X_j, \forall j \in [n' - \gamma n]$$

and values

$$X'_k = X_{n' - \gamma n + k}, \forall k \in [n']$$

Then applying the algorithm gives a stopping time τ with the guarantee

$$\mathbb{E}[X'_\tau] \geq \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\max\{X'_1, \ldots, X'_{n'}\}] \geq \frac{1+\gamma}{2} \cdot \left(1 - \frac{1}{e}\right) \mathbb{E}[\max\{X'_1, \ldots, X'_{n'}\}]$$

as desired. □

We note that we can extend the $1/e$ upper bound on $(0, n)$-stopping rules to also give a parametric upper bound on $(\gamma n, n)$-stopping rules. We again omit the proof, but it can be found in [CDFS21].

Corollary 5.2. *Any* $(\gamma n, n)$*-stopping rule must achieve a competitive ratio of at most* $f(\gamma)$*, where*

$$f(\gamma) = \begin{cases} \dfrac{1+\gamma}{e} & \text{if } \dfrac{\gamma}{1+\gamma} \leq \dfrac{1}{e} \\ -\gamma \cdot \log \dfrac{\gamma}{1+\gamma} & \text{otherwise} \end{cases}$$

6 An improved lower bound with γn samples for $\gamma \leq 1$ in [KNR19]

Following the work in [CDFS21], [KNR19] examines an algorithm that is more careful than simply setting the number of samples and values to be equal, which is able to achieve an improved lower bound for $(\gamma n, n)$-stopping rules. The algorithm works as follows.

Algorithm 3: Parametric Threshold Generation with γn Samples

```
Data: Sequence of i.i.d. random variables $X_1, \ldots, X_n$ sampled from an unknown distribution
      $\mathcal{D}$, with sample access to $\mathcal{D}$
Result: Stopping time $\tau$
$S_1, \ldots, S_{\gamma n} \leftarrow \gamma n$ samples from $\mathcal{D}$
$S \leftarrow \{S_1, \ldots, S_{\gamma n}\}$
$q \leftarrow \max\left\{0, e^{-e^{-\gamma}} - \gamma\right\}$
for $i = 1, \ldots, n$ do
    $S \leftarrow S \cup X_i$
    if $i \leq qn$ then
        Continue
    else
        if $|S| \leq n$ then
            if $X_i = \max S$ then
                return $i$
            end
        else
            $T_i \leftarrow$ random subset of size $n - 1$ from $S$
            if $X_i > \max T_i$ then
                return $i$
            end
        end
    end
end
```

We note a few things about this algorithm. As a high level overview, we first include more samples from the initial values until there are qn samples, where q is a parameter of γ. Then if there are less than $n-1$ samples, the algorithm outputs X_i if it is the maximum so far; otherwise, it uses a random subset of $n-1$ values seen so far as a threshold, as before.

Finally, note that the solution of $e^{-e^{-\gamma}} - \gamma \geq 0$ is at $\gamma \leq r \approx 0.567$. We have the following.

Theorem 6.1. *There exists a* $(\gamma n, n)$*-stopping rule that achieves a competitive ratio of* $f(y)$*, where*

$$f(y) = \begin{cases} e^{-e^{-\gamma}} & \text{if } \gamma \le r \\ \gamma(1 - \ln \gamma - e^{-\gamma}) & \text{otherwise} \end{cases}$$

We omit the proof, which can be found in [KNR19]. Note that this matches the $1/e$ bound for $\gamma = 0$ and $1 - 1/e$ for $\gamma = 1$, but improves on the previous parametric bound in between.

7 Another improved lower bound with γn samples for $\gamma < 1.44$ in [CDF$^+$20]

Following the work in [KNR19], [CDF$^+$20] further improves the bound for γn samples for all $\gamma > 0$. The idea is to vary the amount of samples that are used in the calculation of the threshold rather than mostly using $n - 1$. These are called *maximum-of-random-subset* algorithms, or *MRS* algorithms.

In particular, they consider algorithms that fix a function $f : [n] \to \mathbb{N}$ such that value X_i is accepted if it is greater than $f(i)$ samples and values seen so far, which we previously showed was equivalent to $f(i)$ i.i.d. samples. Note that letting $f(i) = n - 1$ for all i gives exactly the previous bound of $1 - 1/e$ for $n - 1$ samples.

The analysis is straightforward, but very computationally tedious, so we omit the details here. The idea is simply to use numerical methods to solve for the optimal values of f. Noteworthy results include the following.

Theorem 7.1. *For* $0 \le \gamma \le 1/(e-1)$*, there exists a* $(\gamma n, n)$*-stopping rule that achieves a competitive ratio of* $(1 + \gamma)/e$*. By the parametric upper bound given in [CDFS21], this is tight.*

Theorem 7.2. *There exists an* (n, n)*-stopping rule that achieves a competitive ratio of* $\approx 0.6489 \ge 1 - 1/e$*.*

Theorem 7.3. *The best MRS algorithm gives a* $(\approx 1.44n, n)$*-stopping rule that achieves a competitive ratio of* ≈ 0.6534*.*

8 A tight lower bound with $O(n)$ samples in [RWW19]

An algorithm that achieves the upper bound of the optimal competitive ratio of 0.745 with $O(n)$ samples is defined and guaranteed by [RWW19], which ensures that the bound is tight for the i.i.d. prophet inequality with linear samples. This result resolves an open problem from [CDFS21], where the authors of [CDFS21] found that the optimal competitive ratio of $\alpha - \varepsilon$ where $\alpha \approx 0.745$ is achievable with at least $\Omega(n)$ samples and is therefore impossible for $O(n)$ samples. Rubinstein et al. disprove this by defining the Samples-CFHOV algorithm in [RWW19], which modifies the algorithm used in [CFH$^+$17] and [CDFS21] denoted by Explicit-CFHOV, and they prove that for $O(n/\varepsilon^6)$ samples, the Samples-CFHOV algorithm achieves at least a $(1 - O(\varepsilon))$ fraction of the expected reward of Explicit-CFHOV.

The Explicit-CFHOV algorithm sets a probability p_i independent of $\mathcal{D}$ for each $i \in [n]$, and sets a threshold σ_i for stopping at X_i, which is exceeded with probability exactly p_i and is identical to stopping at X_i if and only if $F_D(X_i) > 1 - p_i$.

Theorem 8.1 ([CFH$^+$17], [CDFS21])**.** *The Explicit-CFHOV algorithm has competitive ratio $\alpha - \varepsilon$ in the i.i.d. setting.*

The Explicit-CFHOV algorithm depends on explicit access to $\mathcal{D}$ to be able to exactly compute $F_D(X_i)$. However, if instead we have m i.i.d. samples from $\mathcal{D}$, we must derive another algorithm that will allow us to estimate the distribution $\mathcal{D}$. While [CDFS21] showed that $m = O(n^2)$ samples can estimate $\mathcal{D}$ sufficiently well, [RWW19] observed that only $m = O(n)$ samples suffice.

The authors define the Samples-CFHOV Algorithm as follows:

1. As a function only of n, and independently of $\mathcal{D}$, define monotone increasing probabilities $0 \leq p_1 \leq \cdots p_n \leq 1$.
2. Round down each p_i to the nearest integer power of $(1+\varepsilon)$. Denote the rounded value by $\lfloor p_i \rfloor \in \{(1+\varepsilon)^{-1}, (1+\varepsilon)^{-2}, \ldots\}$.
3. Set $\tilde{p}_i := \lfloor p_i \rfloor / (1+\varepsilon)$.
4. Using our m samples, let τ_i denote the value of the $(\tilde{p}_i \cdot m)$-th highest sample.
5. Stop at X_i if and only if $X_i > \tau_i$.

Thus, Samples-CFHOV provides an estimate τ_i of the σ_i used in Explicit-CFHOV based on the m samples, and Samples-CFHOV attempts to overestimate σ_i so that it is unlikely that Samples-CFHOV will ever choose to stop at a number that Explicit-CFHOV would not stop at.

This algorithm is similar to the one described in [CDFS21] in that they both attempt to set thresholds τ_i such that $F_D(\tau_i) \approx 1 - p_i$, but this algorithm targets a multiplicative $(1-\varepsilon)$ approximation for each threshold while the $O(n^2)$ algorithm targets an additive $1/n$ approximation. In other words, this algorithm seeks a weaker guarantee such that $|F_D(\tau_i) - p_i| \leq O(\varepsilon p_i)$ while the other algorithm guarantees $|F_D(\tau_i) - p_i| \leq 1/n$. In seeking this weaker guarantee that still provides "good" thresholds, we are able to use much less samples on the order of $O(n)$ to solve the i.i.d. prophet inequality.

Theorem 8.2. *With $O(n/\varepsilon^6)$ samples, Samples-CFHOV achieves a competitive ratio of $\alpha - O(\varepsilon)$.*

The authors prove Theorem 8.2 by showing that the expected value of Samples-CFHOV is at least a $1 - O(\varepsilon)$ fraction of that of Explicit-CFHOV. To show this, they make two claims: 1) $O(n)$ samples suffices to determine "good" thresholds with high probability, and 2) these "good" thresholds yield a good approximation of the Explicit-CFHOV thresholds.

Lemma 8.3 (Thresholds are "good" with high probability)**.** *With $\delta = \varepsilon^2/n$ and $m = 12\ln(1/\varepsilon)/(\varepsilon^3\delta) = O(n/\varepsilon^6)$ samples, with probability at least $1-\varepsilon$, we have simultaneously for every i*

$$\frac{p_i}{(1+\varepsilon)^3} \leq \mathbb{P}_{x\sim\mathcal{D}}[x > \tau_i] \leq p_i. \tag{1}$$

Note that this equation does not reference the values of the actual elements X_i, it simply makes a claim about the thresholds τ_i, and thus the probability $1-\varepsilon$ is taken only over the randomness in drawing the samples. A set of thresholds are "good" if they satisfy Equation (1).

Proof. Recall that $\tau_i = \frac{\lfloor p_i \rfloor \cdot m}{(1+\varepsilon)}$. Define L_i to be the random variable such that $\mathbb{P}_{x\sim\mathcal{D}}[x > L_i] = \lfloor p_i \rfloor$ and define H_i to be the random variable such that $\mathbb{P}_{x\sim\mathcal{D}}[x > H_i] = (1+\varepsilon)^{-2}\lfloor p_i \rfloor$. Then, we can prove (1) by showing that $L_i < \tau_i < H_i$ for all i with high probability.

Specifically, we expect to see $\lfloor p_i \rfloor m$ samples greater than L_i, and we say that $\lfloor p_i \rfloor$ is "bad" if the number of samples greater than L_i is not in the range $[(1+\varepsilon)^{-1}(\lfloor p_i \rfloor m), (1+\varepsilon)(\lfloor p_i \rfloor m)]$. When neither $\lfloor p_i \rfloor$ nor $(1+\varepsilon)^{-2}\lfloor p_i \rfloor$ is bad, then we indeed have $L_i < \tau_i < H_i$. Then, we only need to bound the probability that any particular p is bad.

We can use multiplicative Chernoff bounds to get that the probability that a particular p is bad is upper bounded by

$$\mathbb{P}[p \text{ is bad}] < e^{-\varepsilon^2 pm/3}.$$

and then take union bound over all $p \in \{(1+\varepsilon)^{-1}, (1+\varepsilon)^{-2}, \ldots, \delta\}$ to get the probability that some p is bad is bounded by

$$\sum_{i=0}^{O(\ln(1/\delta)/\varepsilon)} e^{-\varepsilon^2(1-\varepsilon)^{-i}\delta m/3} \leq \sum_{i=0}^{\infty} e^{-\varepsilon^2(1-\varepsilon)^{-i}\delta m/3} \leq \sum_{i=0}^{\infty} e^{-\varepsilon^3 i\delta m/6} \leq e^{-\varepsilon^3\delta m/12}$$

where we start with the a union bound over this $(1+\varepsilon)$ -multiplicative net, and the first inequality simply extends the sum to infinity. The second inequality follows as $(1-\varepsilon)^{-i} \geq \varepsilon i/2$ for all $\varepsilon \in (0,1)$ and $i \geq 0$. The final inequality holds at least when $m \geq 6/(\varepsilon^3\delta)$. Thus, setting $m = 12\ln(1/\varepsilon)/(\varepsilon^3\delta)$ satisfies the claim with probability at least $1-\varepsilon$. □

For the second claim, let t_1 be the stopping time of Explicit-CFHOV, and let t_2 be the stopping time of Samples-CFHOV.

Claim 8.4. *Conditioned on (1) holding for every i, $t_2 \geq t_1$. In other words, Samples-CFHOV stops at an element later than Explicit-CFHOV.*

Proof. This is clearly true since by (1) the threshold used by Samples-CFHOV is greater than or equal to the threshold used by Explicit-CFHOV. Thus, if the algorithms differ at all, it is when Explicit-CFHOV chooses an element but Samples-CFHOV does not. □

Lemma 8.5 ("Good" thresholds are a good approximation)**.** *Conditioned on (1) holding for every i, the following holds for every v:*

$$\mathbb{P}[X_{t_2} > v] \geq \frac{1}{(1+\varepsilon)^3}\mathbb{P}[X_{t_1} > v].$$

In other words, this lemma asserts that when the thresholds are "good", Samples-CFHOV achieves at least a $1/(1+\varepsilon)^3$ fraction of the expected reward of Explicit-CFHOV. This is because the expected reward of Explicit-CFHOV is $\int_0^\infty \mathbb{P}[X_{t_1} > v]dv$ while the expected reward of Samples-CFHOV is

$$\int_0^\infty \mathbb{P}[X_{t_2} > v]dv \geq \int_0^\infty \frac{\mathbb{P}[X_{t_1} > v]}{(1+\varepsilon)^3}dv = \frac{1}{(1+\varepsilon)^3} \cdot \mathbb{E}[\text{Explicit-CFHOV}].$$

Proof. We would like to prove that the claim holds for every $i \in [n]$, in other words, the equivalent inequality:

$$\mathbb{P}[(X_{t_2} > v) \wedge (t_2 = i)] \geq \frac{1}{(1+\varepsilon)^3}\mathbb{P}[(X_{t_1} > v) \wedge (t_1 = i)].$$

The event $(X_{t_b} > v) \wedge (t_b = i)$ for both $b = 1, 2$ happens if and only if the corresponding algorithm (Explicit or Samples) does not stop before i and X_i is larger than both v and the threshold set for i by the algorithm. We show that the inequality holds by proving that even though τ_i is a stricter threshold than σ_i, we are still roughly as likely to accept an X_i exceeding v using τ_i versus σ_i for all v. This is the following claim:

Claim 8.6. *Conditioned on Equation (1) holding for every i, then for every v and i such that $p_i \geq \delta$:*

$$(1+\varepsilon)^3\mathbb{P}[(X_i > v) \wedge (X_i > \tau_i)] \geq \mathbb{P}[(X_i > v) \wedge (X_i > \sigma_i)].$$

Proof. We prove this by showing it holds under the three following cases for v:

1. $v \geq \tau_i$: if $v \geq \tau_i \geq \sigma_i$, we clearly must have

$$\mathbb{P}[(X_i > v) \wedge (X_i > \tau_i)] = Pr[X_i > v] = \mathbb{P}[(X_i > v) \wedge (X_i > \sigma_i)].$$

2. $v \in (\sigma_i, \tau_i)$: this implies that

$$\begin{aligned}\mathbb{P}[(X_i > v) \wedge (X_i > \sigma_i)] &\leq \mathbb{P}[X_i > \sigma_i] \\ &\leq (1+\varepsilon)^3\mathbb{P}[X_i > \tau_i] \\ &= (1+\varepsilon)^3\mathbb{P}[(X_i > v) \wedge (X_i > \tau_i)]\end{aligned}$$

   where the first inequality follows due to $v > \sigma_i$ and the second inequality follows from condition 1, and the third equality is by definition of $v < \tau_i$.

3. $v < \sigma_i$: this implies that

$$\begin{aligned}\mathbb{P}[(X_i > v) \wedge (X_i > \sigma_i)] &= \mathbb{P}[X_i > \sigma_i] \\ &\leq (1+\varepsilon)^3\mathbb{P}[X_i > \tau_i] \\ &= (1+\varepsilon)^3\mathbb{P}[(X_i > \tau_i) \wedge (X_i > v)]\end{aligned}$$

   which simply follows from condition 1.

□

Now, note that $\mathbb{P}[(X_{t_2} > v) \wedge (t_2 = i)] = \mathbb{P}[t_2 \geq i] \cdot \mathbb{P}[(X_i > v) \wedge (X_i > \tau_i)]$ and $\mathbb{P}[(X_{t_1} > v) \wedge (t_1 = i)] = \mathbb{P}[t_1 \geq i] \cdot \mathbb{P}[(X_i > v) \wedge (Xi > \sigma_i)]$. Therefore, we have proven the desired inequality for every $i \in [n]$, as

$$\begin{gathered}(1+\varepsilon)^3\mathbb{P}[(X_i > v) \wedge (X_i > \tau_i)] \geq \mathbb{P}[(X_i > v) \wedge (X_i > \sigma_i)]) \\ (1+\varepsilon)^3\mathbb{P}[(X_i > v) \wedge (X_i > \tau_i)] \cdot \mathbb{P}[t_2 \geq i] \geq \mathbb{P}[(X_i > v) \wedge (X_i > \sigma_i)]) \cdot \mathbb{P}[t_1 \geq i] \\ (1+\varepsilon)^3\mathbb{P}[(X_{t_2} > v) \wedge (t_2 = i)] \geq \mathbb{P}[(X_{t_1} > v) \wedge (t_1 = i)] \\ (1+\varepsilon)^3\mathbb{P}[X_{t_2} > v] \geq \mathbb{P}[X_{t_1} > v]\end{gathered}$$

where we start with Claim 8.6, we apply Claim 8.4, and the rest is by observation of probabilities. This completes the proof. ☐

Proof of Theorem 8.2. Finally, we apply Lemmas 8.3 and 8.5 to prove Theorem 8.2. Lemma 8.3 shows that the thresholds of Samples-CFHOV are good with probability at least $1-\varepsilon$, and Lemma 8.5 shows that whenever thresholds are good, Samples-CFHOV achieves at least a $1/(1+\varepsilon)^3$ fraction of the expected reward of Explicit-CFHOV. All together, Samples-CFHOV achieves at least a $\frac{1-\varepsilon}{(1+\varepsilon)^3}$ fraction of the expected reward of Explicit-CFHOV. ☐

Thus, by Theorem 8.2, because the Samples-CFHOV algorithm achieves at least a $(1-O(\varepsilon))$ fraction of the expected reward of Explicit-CFHOV, this implies that Samples-CFHOV achieves optimal competitive ratio $\alpha-\varepsilon$ for the i.i.d. setting with $O(n)$ samples.

9 Extensions

9.1 Prophet inequalities and streaming algorithms in [CDF+20]

There is a clear extension for prophet inequalities to streaming algorithms, where we are primarily concerned with the space usage that our algorithms use. In the previous sections, we store all of the samples and values that we have seen such that we can randomly choose subsets to create threshold values. Perhaps remarkably, [CDF+20] shows that we can achieve very close bounds even while restricting the space usage:

Theorem 9.1. *Let $\epsilon > 0$. Assume there exists an MRS algorithm with guarantee α using $O(n)$ samples. Further assume that the MRS algorithm is based on a continuous function g with $|\{x \in [0,1] : \exists q \in \mathbb{N} : g(x) = q \cdot \epsilon\}| = O_\epsilon(1)$. Then there exists a streaming algorithm using $O_\epsilon(\log n)$ space with guarantee $\alpha - \epsilon$.*

9.2 A generalization of the prophet inequality problem in [CCES21]

In the previous sections, we saw that the optimal competitive ratio of 0.745 can be guaranteed and is tight for the i.i.d. prophet inequality problem with kn samples from the distribution. Recent work is beginning to extend the single-choice i.i.d. prophet inequality problem to more general versions of the problem which also guarantee the tight optimal competitive ratio of 0.745.

Formally, the authors of [CCES21] study a generic version of the classic single-choice optimal stopping problem with sampling, which they call the p-sample-driven optimal stopping problem (p-DOS). In this setup, a collection of n items is shuffled in uniform random order. Instead of being given k offline samples, the decision maker (DM) initially gets to sample each of the n items independently with probability $p \in [0,1)$ and can observe the relative rankings of these sampled items. We call this set of samples the information set or history set. Then, the DM views the remaining items in sequential order and they must decide whether to stop or continue on each item, as in the classic prophet inequality. Moreover, the DM's reward for stopping at the i-th ranked item is Y_i, and the goal of the DM is to maximize their reward. We may assume the rewards are monotone, i.e. $Y_1 \geq \cdots Y_n$, but we do not assume that they are non-negative. We would like to measure the performance of an algorithm that solves this problem by determining the expected maximum value in the online set over the permutations of the items of the collection.

The authors discuss both the case in which the rewards Y_i are known to the DM and the case in which the rewards Y_i are chosen by an adversary. They used a linear program to describe the problem, and they prove that this LP exactly encodes the optimal algorithm. They further derive a limit LP, through which they can limit the inequalities and ranges that are tight in an optimal solution, leading them to solve certain simple ODEs that provide thresholds t_i which are the times at which the DM should accept an item of rank i or higher among the items they have seen thus far.

In the latter case with an adversary, the authors show that the LP holds with the addition of a stochastic dominance constraint and thus can be solved for the optimal algorithm, which takes the form of a sequence of thresholds t_i. Based on these algorithms, the authors prove a series of quantitative results for different values of p, and specifically their guarantee for p approaches 1 of optimal performance that approaches 0.745, which matches that of the i.i.d. prophet inequality. This implies that there is no loss by considering a more general combinatorial version of the prophet inequality problem without full distributional knowledge.

10 Conclusion

In this paper, we have shown a variety of algorithms and stopping rules used in research on the single-choice i.i.d. prophet inequality problem with sampling, and in particular we discuss each algorithm's performance in terms of the optimal competitive ratio α achieved by the algorithm, i.e. α such that the algorithm guarantees expected value on stopping at a certain element that is at least an α fraction of the best expected value over all the elements. Researchers have been able to continuously improve on a lower bound for the optimal competitive ratio with kn samples, eventually showing that the upper bound of 0.745 for the general prophet inequality problem (where the distribution is known) is achievable with $O(n)$ samples in the sampling prophet inequality problem.

Recent work related to the problem of prophet inequality problem with sampling include applications to other fields, such as streaming algorithms (Section 9.1), and extensions to generalizations of the problem, such as in combinatorial situations (Section 9.2) and the matroid secretary problem with sampling. One such related open problem is the existence of a constant competitive algorithm for the matroid secretary problem, which may be tackled following the algorithms and results of the p-DOS setup of [CCES21]. Additionally, it is worth noting that the prophet inequality problem with sampling has been interesting to researchers because of its real-world applications to algorithmic pricing such as posted-price mechanisms and choice of prices in online advertising auctions. We expect future research to apply knowledge in solving the secretary problem and prophet inequality problem to other related applications, ranging from price mechanisms to streaming algorithms to combinatorial auctions.

References

[CCES19] José Correa, Andrés Cristi, Boris Epstein, and José A. Soto. The two-sided game of googol and sample-based prophet inequalities, 2019.

[CCES21] José Correa, Andrés Cristi, Boris Epstein, and José A. Soto. Sample-driven optimal stopping: from the secretary problem to the i.i.d. prophet inequality, 2021.

[CDF$^+$20] José Correa, Paul Dütting, Felix Fischer, Kevin Schewior, and Bruno Ziliotto. Unknown i.i.d. prophets: Better bounds, streaming algorithms, and a new impossibility, 2020.

[CDFS21] José R. Correa, Paul Dütting, Felix Fischer, and Kevin Schewior. Prophet inequalities for i.i.d. random variables from an unknown distribution, 2021.

[CFH$^+$17] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17 Cambridge, MA, USA, June 26-30, 2017*, pages 169–186, 2017.

[HK82] T. P. Hill and Robert P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability*, 10(2):336–345, 1982.

[Ker86] Robert P Kertz. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis*, 19(1):88–112, 1986.

[KNR19] Haim Kaplan, David Naori, and Danny Raz. Competitive analysis with a sample and the secretary problem, 2019.

[RWW19] Aviad Rubinstein, Jack Z. Wang, and S. Matthew Weinberg. Optimal single-choice prophet inequalities from samples, 2019.