Prophet Secretary*

Hossein Esfandiari† MohammadTaghi Hajiaghayi‡ Vahid Liaghat §
Morteza Monemizadeh¶

July 5, 2015

Abstract

Optimal stopping theory is a powerful tool for analyzing scenarios such as online auctions in which we generally require optimizing an objective function over the space of stopping rules for an allocation process under uncertainty. Perhaps the most classic problems of stopping theory are the prophet inequality problem and the secretary problem. The classical prophet inequality states that by choosing the same threshold $OPT/2$ for every step, one can achieve the tight competitive ratio of 0.5. On the other hand, for the basic secretary problem, the optimal strategy achieves the tight competitive ratio of $1/e \approx 0.36$

In this paper, we introduce *prophet secretary*, a natural combination of the prophet inequality and the secretary problems. An example motivation for our problem is as follows. Consider a seller that has an item to sell on the market to a set of arriving customers. The seller knows the types of customers that may be interested in the item and he has a price distribution for each type: the price offered by a customer of a type is anticipated to be drawn from the corresponding distribution. However, the customers arrive in a random order. Upon the arrival of a customer, the seller makes an irrevocable decision whether to sell the item at the offered price. We address the question of finding a strategy for selling the item at a high price. In particular, we show that by using a single uniform threshold one cannot break the 0.5 barrier of the prophet inequality for the prophet secretary problem. However, we show that

- using n distinct non-adaptive thresholds one can obtain a competitive ratio that goes to $(1-1/e \approx 0.63)$ as n grows; and
- no online algorithm can achieve a competitive ratio better than 0.73.

Our results improve the (asymptotic) approximation guarantee of single-item sequential posted pricing mechanisms from 0.5 to $(1-1/e)$ when the order of agents (customers) is chosen randomly.

We also consider the minimization variants of stopping theory problems and in particular the prophet secretary problem. Interestingly, we show that, even for the simple case in which the input elements are drawn from identical and independent distributions (i.i.d.), there is no constant competitive online algorithm for the minimization variant of the prophet secretary problems. We extend this hardness result to the minimization variants of both the prophet inequality and the secretary problem as well.

*A preliminary version of this paper is to appear in Proceedings of the 23rd European Symposium on Algorithms (ESA 2015), Patras, Greece.

†University of Maryland, College Park, MD, USA. Supported in part by NSF CAREER award 1053605, NSF grant CCF-1161626, ONR YIP award N000141110662, DARPA/AFOSR grant FA9550-12-1-0423. hossein@cs.umd.edu.

‡University of Maryland, College Park, MD, USA. Supported in part by NSF CAREER award 1053605, NSF grant CCF-1161626, ONR YIP award N000141110662, DARPA/AFOSR grant FA9550-12-1-0423. Hajiagha@cs.umd.edu.

§Management Science and Engineering and Institute for Computational and Mathematical Engineering, Stanford University. Supported in part by a Google PhD Fellowship. vliaghat@stanford.edu.

¶Computer Science Institute of Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic. Partially supported by the project 14-10003S of GA ČR. Part of this work was done when the author was at Department of Computer Science, Goethe-Universität Frankfurt, Germany and supported in part by MO 2200/1-1. monemi@iuuk.mff.cuni.cz.

1 Introduction

Optimal stopping theory is a powerful tool for analyzing scenarios in which we generally require optimizing an objective function over the space of stopping rules for an allocation process under uncertainty. One such a scenario is the online auction which is the essence of many modern markets, particularly networked markets where information about goods, agents, and outcomes is revealed over a period of time and the agents must make irrevocable decisions without knowing future information. Combining optimal stopping theory with game theory allows us to model the actions of rational agents applying competing stopping rules in an online market.

Perhaps the most classic problems of stopping theory are the *prophet inequality* and the *secretary problem*. Research investigating the relation between online auction mechanisms and prophet inequalities was initiated by Hajiaghayi, Kleinberg, and Sandholm [17]. They observed that algorithms used in the derivation of prophet inequalities, owing to their monotonicity properties, could be interpreted as truthful online auction mechanisms and that the prophet inequality in turn could be interpreted as the mechanism's approximation guarantee. Later Chawla, Hartline, Malec, and Sivan [11] showed the applications of prophet inequalities in Bayesian optimal mechansim design problems. The connection between the secretary problem and online auction mechanisms has been explored by Hajiaghayi, Kleinberg and Parkes [16] and initiated several follow-up papers (see e.g. [7, 8, 9, 18, 20]).

Prophet Inequality. The classical prophet inequality has been studied in the optimal stopping theory since the 1970s when introduced by Krengel and Sucheston [19, 22, 23] and more recently in computer science Hajiaghayi, Kleinberg and Sandholm [17]. In the prophet inequality setting, given (not necessarily identical) distributions $\{D_1, \ldots, D_n\}$, an online sequence of values $X_1, \cdots, X_n$ where X_i is drawn from D_i, an onlooker has to choose one item from the succession of the values, where X_k is revealed at step k. The onlooker can choose a value only at the time of arrival. The onlooker's goal is to maximize her revenue. The inequality has been interpreted as meaning that a prophet with complete foresight has only a bounded advantage over an onlooker who observes the random variables one by one, and this explains the name *prophet inequality*.

An algorithm for the prophet inequality problem can be described by setting a threshold for every step: we stop at the first step that the arriving value is higher than the threshold of that step. The classical prophet inequality states that by choosing the same threshold $\mathsf{OPT}/2$ for every step, one achieves the competitive ratio of $1/2$. Here the optimal solution OPT is defined as $E[\max X_i]$. Naturally, the first question is whether one can beat $1/2$. Unfortunately, this is not possible: let $q = \frac{1}{\epsilon}$, and $q' = 0$. The first value X_1 is always 1. The second value X_2 is either q with probability ϵ or q' with probability $1 - \epsilon$. Observe that the expected revenue of any (randomized) online algorithm is $\max(1, \epsilon(\frac{1}{\epsilon})) = 1$. However the prophet, i.e., the optimal offline solution would choose q' if it arrives, and he would choose the first value otherwise. Hence, the optimal offline revenue is $(1 - \epsilon) \times 1 + \epsilon(\frac{1}{\epsilon}) \approx 2$. Therefore we cannot hope to break the $1/2$ barrier using any online algorithm.

Secretary Problem. Imagine that you manage a company, and you want to hire a secretary from a pool of n applicants. You are very keen on hiring only the best and brightest. Unfortunately, you cannot tell how good a secretary is until you interview him, and you must make an irrevocable decision whether or not to make an offer at the time of the interview. The problem is to design a strategy which maximizes the probability of hiring the most qualified secretary. It is well-known since 1963 by Dynkin in [12] that the optimal policy is to interview the first $t - 1$ applicants, then hire the next one whose quality exceeds that of the first $t - 1$ applicants, where t is defined by $\sum_{j=t+1}^{n} \frac{1}{j-1} \leq 1 < \sum_{j=t}^{n} \frac{1}{j-1}$. As $n \to \infty$, the probability of hiring the best applicant approaches $1/e \approx 0.36$, as does the ratio t/n. Note that a solution to

the secretary problem immediately yields an algorithm for a slightly different objective function optimizing the expected value of the chosen element. Subsequent papers have extended the problem by varying the objective function, varying the information available to the decision-maker, and so on, see e.g., [1, 15, 26, 27].

1.1 Further Related Work

Prophet Inequality. The first generalization of the basic prophet inequality introduced by Krengel and Sucheston [19, 22, 23] is *the multiple-choices* prophet inequality [6] in which both the onlooker and the prophet have $k > 1$ choices. Currently, the best algorithm for this setting is due to Alaei [2], who gave an online algorithm with $(1 - \frac{1}{\sqrt{k+3}})$-competitive ratio for k-choice optimal stopping. Besides this, we have two generalizations for the (multiple-choices) prophet inequality that are *matroid* prophet inequality [21] and *matching* prophet inequality [3].

In the matroid prophet inequality, we are given a matroid whose elements have random weights sampled independently from (not necessarily identical) probability distributions on $\mathbb{R}^+$. We then run an online algorithm with knowledge of the matroid structure and of the distribution of each element's weight. The online algorithm must then choose irrevocably an independent subset of the matroid by observing the sampled value of each element (in a fixed, prespecified order). The online algorithm's payoff is defined to be the sum of the weights of the selected elements. Kleinberg and Weinberg [21] show that for every matroid, there is an online algorithm whose expected payoff is at least half of the expected weight of the maximum-weight basis. Observe that the original prophet inequality introduced by Krengel and Sucheston [19, 22, 23] corresponds to the special case of rank-one matroids.

The matching prophet inequality is due to Alaei, Hajiaghayi, and Liaghat [5, 3, 4]. Indeed, they study the problem of online prophet-inequality matching in bipartite graphs. There is a static set of bidders and an online stream of items. The interest of bidders in items is represented by a weighted bipartite graph. Each bidder has a capacity, i.e., an upper bound on the number of items that can be allocated to him. The weight of a matching is the total weight of edges matched to the bidders. Upon the arrival of an item, the online algorithm should either allocate it to a bidder or discard it. The objective is to maximize the weight of the resulting matching. Here we assume we know the distribution of the incoming items in advance and we may assume that the tth item is drawn from distribution $\mathcal{D}_t$. They generalize the $\frac{1}{2}$-competitive ratio of Krengel and Sucheston [22] by presenting an algorithm with an approximation ratio of $1 - \frac{1}{\sqrt{k+3}}$ where k is the minimum capacity. Oberve that the classical prophet inequality is a special case of this model where we have only one bidder with capacity one, i.e., $k = 1$ for which they get the same $\frac{1}{2}$-competitive ratio.

Secretary Problem. The first generalization of the basic secretary problem [12] is the *multiple-choice secretary problem* [8] (see a survey by Babaioff et al. [8])) in which the interviewer is allowed to hire up to $k \geq 1$ applicants in order to maximize performance of the secretarial group based on their overlapping skills (or the joint utility of selected items in a more general setting). More formally, assuming applicants of a set $S = \{a_1, a_2, \cdots, a_n\}$ (applicant pool) arriving in a uniformly random order, the goal is to select a set of at most k applicants in order to maximize a non-negative profit function $f : 2^S \mapsto \mathbb{R}^{\geq 0}$. For example, when $f(T)$ is the maximum individual value [13, 14], or when $f(T)$ is the sum of the individual values in T [20], the problem has been considered thoroughly in the literature. Beside this, two generalizations for the (multiple-choices) secretary problem are *submodular secretary* [10] and *matroid secretary* [9].

The submodular secretary problem is introduced by Bateni, Hajiaghayi, and Zadimoghaddam [10]. Indeed, both of the maximum individual value [13, 14] and the sum of the individual values [20] aforementioned are special monotone non-negative submodular functions. Bateni, Hajiaghayi, and Zadimoghaddam [10] give an online algorithm with $(\frac{7}{1-1/e})$-competitive ratio for the submodular secretary problem. We

should mention that there are more recent results with better constant competitive ratio (See for example the references in [10]).

In the matroid secretary problem considered by Babaioff et al. [9], we are given a matroid with a ground set $\mathcal{U}$ of elements and a collection of independent (feasible) subsets $\mathcal{I} \subseteq 2^{\mathcal{U}}$ describing the sets of elements which can be simultaneously accepted. The goal is to design online algorithms in which the structure of $\mathcal{U}$ and $\mathcal{I}$ is known at the outset (assume we have an oracle to answer whether a subset of $\mathcal{U}$ belongs to $\mathcal{I}$ or not), while the elements and their values are revealed one at a time in a random order. As each element is presented, the algorithm must make an irrevocable decision to select or reject it such that the set of selected elements belongs to $\mathcal{I}$ at all times. Babaioff et al. [9] present an $O(\log r)$-competitive algorithm for general matroids, where r is the rank of the matroid. However, they leave as a main open question the existence of constant-competitive algorithms for general matroids. On the other hand, there are several follow-up works including the recent FOCS 2014 paper due to Lachish [24] which gives $O(\log \log \text{rank})$-competitive algorithm.

2 Our Contributions

In this paper, we introduce *prophet secretary* as a natural combination of the prophet inequality problem and the secretary problem with applications to the Bayesian optimal mechanism design. Consider a seller that has an item to sell on the market to a set of arriving customers. The seller knows the types of customers that may be interested in the item and he has a price distribution for each type: the price offered by a customer of a type is anticipated to be drawn from the corresponding distribution. However, the customers arrive in a random order. Upon the arrival of a customer, the seller makes an irrevocable decision to whether sell the item at the offered price. We address the question of maximizing the seller's gain.

More formally, in the prophet secretary problem we are given a set $\{D_1, \ldots, D_n\}$ of (not necessarily identical) distributions. A number X_i is drawn from each distribution D_i and then, after applying a random permutation $\pi_1, \ldots, \pi_n$, the numbers are given to us in an online fashion, i.e., at step k, π_k and X_{π_k} are revealed. We are allowed to choose only one number, which can be done only upon receiving that number. The goal is to maximize the expectation of the chosen value, compared to the expectation of the optimum offline solution that knows the drawn values in advance (i.e., $\mathsf{OPT} = \mathrm{E}\left[\max_i X_i\right]$). For the ease of notation, in what follows the index i iterates over the distributions while the index k iterates over the arrival steps.

An algorithm for the prophet secretary problem can be described by a sequence of (possibly adaptive) thresholds $\langle \tau_1, \ldots, \tau_n \rangle$: we stop at the first step k that $X_{\pi_k} \geq \tau_k$. In particular, if the thresholds are non-adaptive, meaning that they are decided in advance, the following is a generic description of an algorithm. The competitive ratio of the following algorithm is defined as $\frac{\mathrm{E}[Y]}{\mathsf{OPT}}$.

Algorithm Prophet Secretary
Input: A set of distributions $\{D_1, \ldots, D_n\}$; a randomly permuted stream of numbers $(X_{\pi_1}, \ldots, X_{\pi_n})$ drawn from the corresponding distributions.
Output: A number Y.

(1) Let $\langle \tau_1, \ldots, \tau_n \rangle$ be a sequence of thresholds.

(2) For $k \leftarrow 1$ to n

  (a) If $X_{\pi_k} \geq \tau_k$ then let $Y = X_{\pi_k}$ and exit the For loop.

(3) Output Y as the solution.

Recall that when the arrival order is adversarial, the classical prophet inequality states that by choosing the same threshold $\mathsf{OPT}/2$ for every step, one achieves the tight competitive ratio of $1/2$. On the other hand, for the basic secretary problem where the distributions are not known, the optimal strategy is to let $\tau_1 = \cdots = \tau_{\frac{n}{e}} = \infty$ and $\tau_{\frac{n}{e}+1} = \cdots = \tau_n = \max(X_{\pi_1}, \ldots, X_{\pi_{\frac{n}{e}}})$. This leads to the optimal competitive ratio of $\frac{1}{e} \simeq 0.36$. Hence, our goal in the prophet secretary problem is to beat the $1/2$ barrier.

We would like to mention that in an extension of this problem, in which the seller has B identical items to sell, it is indeed easier to track the optimal solution since we have multiple choices. In fact, an algorithm similar to that of [3, 4] can guarantee a competitive ratio of $1 - \frac{1}{\sqrt{B+3}}$ which goes to one as B grows.

We first show that unlike the prophet inequality setting, one cannot obtain the optimal competitive ratio by using a single uniform threshold. Indeed, in Section 4 we show that $1/2$ is the best competitive ratio one can achieve with uniform thresholds. To beat the $\frac{1}{2}$ barrier, as a warm up we first show in Section 5 that by using two thresholds one can achieve the competitive ratio of $5/9 \simeq 0.55$. This can be achieved by choosing the threshold $\frac{5}{9} \cdot \mathsf{OPT}$ for the first half of the steps and then decreasing the threshold to $\frac{\mathsf{OPT}}{3}$ for the second half of the steps. Later in Section 6, we show that by setting n distinct thresholds one can obtain the $(1 - 1/e \approx 0.63)$-competitive ratio for the prophet secretary problem.

Theorem 1 *Let $\langle \tau_1, \ldots, \tau_n \rangle$ be a non-increasing sequence of n thresholds, such that (i) $\tau_k = \alpha_k \cdot \mathsf{OPT}$ for every $k \in [n]$; (ii) $\alpha_n = \frac{1}{n+1}$; and (iii) $\alpha_k = \frac{n\alpha_{k+1}+1}{n+1}$ for $k \in [n-1]$. The competitive ratio of Algorithm* Prophet Secretary *invoked with thresholds τ_k's is at least α_1. When n goes to infinity, α_1 converges to $1 - 1/e \approx 0.63$.*

Remark 2 *We should mention that Yan in [28] establishes a 1 - 1/e approximation when the designer is allowed to choose the order of arrival. Thus, Theorem 1 can be viewed as improving that result by showing that a random arrival order is sufficient to obtain the same approximation.*

The crux of the analysis of our algorithm is to compute the probability of picking a value x at a step of the algorithm with respect to the threshold factors α_k's. Indeed one source of difficulty arises from the fundamental dependency between the steps: for any step k, the fact that the algorithm has not stopped in the previous steps leads to various restrictions on what we expect to see at the step k. For example, consider the scenario that D_1 is 1 with probability one and D_2 is either 2 or 0 with equal probabilities. Now if the algorithm chooses $\tau_1 = 1$, then it would never happen that the algorithm reaches step two and receives a number drawn from D_2! That would mean we have received a value from D_1 at the first step which is a contradiction since we would have picked that number. In fact, the optimal strategy for this example is to shoot for D_2! We set $\tau_1 = 2$ so that we can ignore the first value in the event that it is drawn from D_1. Then we set $\tau_2 = 1$ so that we can always pick the second value. Therefore in expectation we get $5/4$ which is slightly less than $\mathsf{OPT} = 6/4$.

To handle the dependencies between the steps, we first distinguish between the events for $k \in [n]$ that we pick a value between τ_{k+1} and τ_k. We show that the expected value we pick at such events is indeed highly dependent on $\theta(k)$, the probability of passing the first k elements. We then use this observation to analyze competitive ratio with respect to $\theta(k)$'s and the thresholds factors α_k's. We finally show that the competitive ratio is indeed maximized by choosing the threshold factors described in Theorem 1. In Section 5, we first prove the theorem for the simple case of $n = 2$. This enables us to demonstrate our techniques without going into the more complicated dependencies for general n. We then present the full proof of Theorem 1 in Section 6. We would like to emphasize that our algorithm only needs to know the value of OPT, thus requiring only a weak access to the distributions themselves.

As mentioned before, Bayesian optimal mechanism design problems provide a compelling application of prophet inequalities in economics. In such a Bayesian market, we have a set of n agents with private types sampled from (not necessary identical) known distributions. Upon receiving the reported types, a seller has to allocate resources and charge prices to the agents. The goal is to maximize the seller's revenue in equilibrium. Chawla et al. [11] pioneered the study the approximability of a special class of such mechanisms, *sequential posted pricing* (SPM): the seller makes a sequence of take-it-or-leave-it offers to agents, offering an item for a specific price. They show although simple, SPMs approximate the optimal revenue in many different settings. Therefore prophet inequalities directly translate to approximation factors for the seller's revenue in these settings through standard machineries. Indeed one can analyze the so-called *virtual values* of winning bids introduced by Roger Myerson [25], to prove via prophet inequalities that the expected virtual value obtained by the SPM mechanism approximates an offline optimum that is with respect to the exact types. Chawla et al. [11] provide a type of prophet inequality in which one can choose the ordering of agents. They show that under matroid feasibility constraints, one can achieve a competitive ratio of 0.5 in this model, and no algorithm can achieve a ratio better 0.8. Kleinberg and Weinberg [21] later improved there result by giving an algorithm with the tight competitive ratio of 0.5 for an adversarial ordering. Our result can be seen as improving their approximation guarantees to 0.63 for the case of single-item SPMs when the order of agents are chosen randomly.

On the other hand, from the negative side the following theorem shows that no online algorithm can achieve a competitive ratio better than 0.73. The proof is given in Section 7.

Theorem 3 *For any arbitrary small positive number ϵ, there is no online algorithm for the prophet secretary problem with competitive ratio $\frac{11}{15} + \epsilon \approx 0.73 + \epsilon$.*

We also consider the minimization variants of the prophet inequality problem, the prophet secretary problem and the secretary problem. In the minimization variant, we need to select one element of the input and we aim to minimize the expected value of the selected element. In particular, we show that, even for the simple case in which numbers are drawn from *identical and independent distributions (i.i.d.)*, there is no $\frac{(1.11)^n}{6}$ competitive online algorithm for the minimization variants of the prophet inequality and prophet secretary problems.

Theorem 4 *The competitive ratio of any online algorithm for the minimization prophet inequality with n identical and independent distributions is bounded by $\frac{(1.11)^n}{6}$. This bound holds for the minimization prophet secretary problem as well.*

Furthermore, we empower the online algorithm and assume that the online algorithm can withdraw and change its decision once. Indeed, for the minimization variants of all, prophet secretary problem, secretary problem and prophet inequality, we show that there is no C competitive algorithm, where C is an arbitrary large number. The proof of Theorem 4 and this last result are given in Section 7.2.

3 Preliminaries

We first define some notation. For every $k \in [n]$, let z_k denote the random variable that shows the value we pick at the kth step. Observe that for a fixed sequence of drawn values and a fixed permutation, at most one of z_k's is non-zero since we only pick one number. Let z denote the value chosen by the algorithm. By definition, $z = \sum_{k=1}^{n} z_k$. In fact, since all but one of z_k's are zero, we have the following proposition. We note that since the thresholds are deterministic, the randomness comes from the permutation π and the distributions.

Proposition 5 $\Pr[z \geq x] = \sum_{k \in [n]} \Pr[z_k \geq x]$.

For every $k \in [n]$, let $\theta(k)$ denote the probability that Algorithm Prophet Secretary does not choose a value from the first k steps. For every $i \in [n]$ and $k \in [n-1]$, let $q_{-i}(k)$ denote the probability that the following two events concurrently happen:

(i) Algorithm Prophet Secretary does not choose a value from the first k elements.

(ii) None of the first k values are drawn from D_i.

Proposition 6 *If the thresholds of Algorithm* Prophet Secretary *are non-increasing, then for every* $i \in [n]$ *and* $k \in [n-1]$*, we have* $\theta(k+1) \leq q_{-i}(k)$.

Proof : In what follows let $i \in [n]$ be a fixed value. The claim is in fact very intuitive: $q_{-i}(k)$ is the probability of the event that the algorithm passes k values chosen from all distributions but D_i. On the other hand, $\theta(k+1)$ corresponds to the event that the algorithm passes $k+1$ values chosen from all distributions. Intuitively, in the latter we have surely passed k values chosen from all but D_i. Therefore $\theta(k+1)$ cannot be more than $q_{-i}(k)$.

Formalizing the intuition above, however, requires an exact formulation of the probabilities. For a permutation s of size k of $[n]$, let $s(j)$, for $j \in [k]$, denote the number at position j of s. For $k \in [n]$, let $S(k)$ denote the set of permutations of size k of $[n]$. Let $S_{-i}(k)$ denote the set of permutations of size k of $[n] \backslash \{i\}$. Observe that

$$|S(k)| = \frac{n!}{(n-k)!} \qquad \text{and} \qquad |S_{-i}(k)| = \frac{(n-1)!}{(n-1-k)!}$$

In particular, we note that $|S(k+1)| = n|S_{-i}(k)|$. We can now write down the exact formula for $q_{-i}(k)$ and $\theta(k+1)$.

$$\theta(k+1) = \frac{1}{|S(k+1)|} \sum_{s \in S(k+1)} \prod_{j \in [k+1]} \Pr\left[X_{s(j)} < \tau_j\right] \tag{1}$$

$$q_{-i}(k) = \frac{1}{|S_{-i}(k)|} \sum_{s \in S_{-i}(k)} \prod_{j \in [k]} \Pr\left[X_{s(j)} < \tau_j\right] \tag{2}$$

We now show that $\theta(k+1)$ can be written down as a convex combination of terms that are less than $q_{-i}(k)$. For every $\ell \in [k]$, let $S(k+1, \ell) = \{s \in S(k+1) | s(\ell) = i\}$. We have

$$\begin{aligned}
\sum_{s:S(k+1,\ell)} \prod_{j \in [k+1]} \Pr\left[X_{s(j)} \leq \tau_j\right] &\leq \sum_{s:S(k+1,\ell)} \prod_{j \in [k+1] \backslash \{\ell\}} \Pr\left[X_{s(j)} \leq \tau_j\right] \\
&\leq \sum_{s:S(k+1,\ell)} \prod_{j \in [k] \backslash \{\ell\}} \Pr\left[X_{s(j)} \leq \tau_j\right] \times \Pr\left[X_{s(k+1)} \leq \tau_\ell\right] \\
&= \sum_{s:S_{-i}(k)} \prod_{j \in [k]} \Pr\left[X_{s(j)} \leq \tau_j\right] \quad ,
\end{aligned}$$

where the first inequality is since $\Pr[X_i < \tau_\ell] \leq 1$ and the second inequality is because $\Pr\left[X_{s(k+1)} \leq \tau_{k+1}\right] \leq \Pr\left[X_{s(k+1)} \leq \tau_\ell\right]$ and since $\tau_{k+1} \leq \tau_\ell$.

Thus, using Equation 2 we have

$$\sum_{s:S(k+1,\ell)} \prod_{j\in[k+1]} \Pr\left[X_{s(j)} \le \tau_j\right] = \sum_{s:S_{-i}(k)} \prod_{j\in[k]} \Pr\left[X_{s(j)} \le \tau_j\right] = |S_{-i}(k)| q_{-i}(k) \ . \tag{3}$$

Equation 3 establishes the relation to $q_{-i}(k)$ for members of $S(k+1)$ that contain i in one of the first k positions.

Let $\overline{S(k+1)} = S(k+1) \setminus \bigcup_{\ell\in[k]} S(k+1,\ell)$. We then have

$$\begin{aligned}
\sum_{s:\overline{S(k+1)}} \prod_{j\in[k+1]\setminus\{\ell\}} \Pr\left[X_{s(j)} \le \tau_j\right] &= \sum_{s:\overline{S(k+1)}} \prod_{j\in[k]\setminus\{\ell\}} \Pr\left[X_{s(j)} \le \tau_j\right] \times \Pr\left[X_{s(k+1)} \le \tau_{k+1}\right] \\
&\le (n-k) \sum_{s:S_{-i}(k+1)} \prod_{j\in[k]\setminus\{\ell\}} \Pr\left[X_{s(j)} \le \tau_j\right] \\
&= (n-k)|S_{-i}(k)| q_{-i}(k) \ ,
\end{aligned}$$

where the inequality is since $\Pr\left[X_{s(k+1)} \le \tau_{k+1}\right] \le 1$.

Finally using this last equation and Equations 1 and 3 and since $|S(k+1)| = n|S_{-i}(k)|$ we obtain

$$\begin{aligned}
\theta(k+1) = \quad & \frac{1}{|S(k+1)|} \sum_{s\in S(k+1)} \prod_{j\in[k+1]} \Pr\left[X_{s(j)} < \tau_j\right] \\
& \frac{1}{|S(k+1)|} \left(\sum_{\ell\in[k]} \sum_{s\in S(k+1,\ell)} \prod_{j\in[k+1]} \Pr\left[X_{s(j)} < \tau_j\right] + \sum_{s\in\overline{S(k+1)}} \prod_{j\in[k+1]} \Pr\left[X_{s(j)} < \tau_j\right] \right) \\
& \le \frac{1}{|S(k+1)|} \left(k|S_{-i}(k)| q_{-i}(k) + (n-k)|S_{-i}(k)| q_{-i}(k)\right) \\
& = q_{-i}(k) \ .
\end{aligned}$$

□

4 One Threshold Cannot Break $\frac{1}{2}$ Barrier for Prophet Secretary

To illustrate that considering at least 2 thresholds is necessary to beat $\frac{1}{2}$ barrier for the prophet secretary problem, we first give an example that shows achieving better than $\frac{1}{2}$-competitive ratio for any online algorithm that uses only one threshold for the prophet secretary problem is not possible.

Theorem 7 *There is no online algorithm for the prophet secretary problem that uses one threshold and can achieve competitive ratio better than* $0.5 + \frac{1}{2n}$.

Proof : Suppose we have $n+1$ distributions where the first n distributions always gives $\frac{1}{1-1/n}$ and the $(n+1)^{\text{th}}$ distribution gives n with probability $\frac{1}{n}$ and gives 0 with probability $1-\frac{1}{n}$. Therefore, with probability $\frac{1}{n}$, the maximum is n, and with probability $1-\frac{1}{n}$, the maximum is $\frac{1}{1-1/n}$. Thus, the expected outcome of the offline optimum algorithm is $\frac{1}{n} \times n + (1-\frac{1}{n}) \times \frac{1}{1-1/n} = 2$.

Now, suppose we have an online algorithm that uses one threshold, say T for a number T, that is the online algorithm accepts the first number greater or equal to a threshold T. We consider two cases for T.

The first case is if $T > \frac{1}{1-1/n}$ for which the algorithm does not accept $\frac{1}{1-1/n}$ and thus, the expected outcome of such an algorithm is $\frac{1}{n} \times n = 1$.

The second case is if $T \leq \frac{1}{1-1/n}$. Observe that, with probability $\frac{n}{n+1}$, the first number is $\frac{1}{1-1/n}$ and the online algorithm accepts it. And, with probability $\frac{1}{n+1}$, the distribution that gives n with probability $\frac{1}{n}$ will be the first and the outcome of the algorithm is 1. Thus, the expected outcome of the online algorithm is

$$\frac{n}{n+1} \times \frac{1}{1-1/n} + \frac{1}{n+1} \times 1 = \frac{n^2}{n^2-1} + \frac{1}{n+1} \leq 1 + \frac{1}{n} \ .$$

Therefore, the competitive ratio of the online algorithms that uses only one threshold is bounded by $\frac{1+1/n}{2} = 0.5 + \frac{1}{2n}$. □

5 Two Thresholds Breaks $\frac{1}{2}$ Barrier

Since using one threshold is hopeless, we now try using two thresholds. More formally, for the first half of steps, we use a certain threshold, and then we use a different threshold for the rest of steps. We note that similar to the one-threshold algorithm, both thresholds should be proportional to OPT. Furthermore, at the beginning we should be optimistic and try to have a higher threshold, but if we cannot pick a value in the first half, we may need to lower the bar! We show that by using two thresholds one can indeed achieve the competitive ratio of $\frac{5}{9} \simeq 0.55$. In fact, this improvement beyond $1/2$ happens even at $n = 2$. Thus as a warm up before analyzing the main algorithm with n thresholds, we focus on the case of $n = 2$.

Let $\tau_1 = \alpha_1 \mathsf{OPT}$ and $\tau_2 = \alpha_2 \mathsf{OPT}$ for some $1 \geq \alpha_1 \geq \alpha_2 \geq 0$ to be optimized later. Recall that z_1 and z_2 are the random variables showing the values picked up by the algorithm at step one and two, respectively. We are interested in comparing $E[z]$ with OPT. By Proposition 5 we have

$$E[z] = \int_0^\infty \Pr[z \geq x]\, dx = \int_0^\infty \Pr[z_1 \geq x]\, dx + \int_0^\infty \Pr[z_2 \geq x]\, dx \ .$$

Observe that z_1 (resp. z_2) is either zero or has a value more than τ_1 (resp. τ_2). In fact, since $\tau_1 \geq \tau_2$, z is either zero or has a value more than τ_2. Recall that $\theta(1)$ is the probability of $z_1 = 0$ while $\theta(2)$ is the probability of $z_1 = z_2 = 0$. This observation leads to the following simplification:

$$\begin{aligned}
E[z] &= \int_0^{\tau_2} \Pr[z_1 \geq x]\, dx + \int_{\tau_2}^{\tau_1} \Pr[z_1 \geq x]\, dx + \int_{\tau_1}^{\infty} \Pr[z_1 \geq x]\, dx + \int_0^{\tau_2} \Pr[z_2 \geq x]\, dx + \int_{\tau_2}^{\infty} \Pr[z_2 \geq x]\, dx \\
&= \int_0^{\tau_2} \Pr[z \geq x]\, dx + \int_{\tau_2}^{\tau_1} \Pr[z_1 \geq x]\, dx \quad + \int_{\tau_1}^{\infty} \Pr[z_1 \geq x]\, dx + \int_{\tau_2}^{\infty} \Pr[z_2 \geq x]\, dx \\
&= \tau_2(1-\theta(2)) + (\tau_1 - \tau_2)(1-\theta(1)) + \int_{\tau_1}^{\infty} \Pr[z_1 \geq x]\, dx + \int_{\tau_2}^{\infty} \Pr[z_2 \geq x]\, dx \ .
\end{aligned}$$

Let us first focus on $\Pr[z_1 \geq x]$. The first value may come from any of the two distributions, thus we have

$$\Pr[z_1 \geq x] = \frac{1}{2} \Pr[X_1 \geq x] + \frac{1}{2} \Pr[X_2 \geq x] \ .$$

On the other hand, z_2 is non-zero only if we do not pick anything at the first step. For $i \in \{1, 2\}$, we pick a value of at least x drawn from D_i at step two, if and only if: (i) the value drawn from D_i is at least x; and (ii) our algorithm does not pick a value from the previous step which is drawn from the other distribution.

By definitions, the former happens with probability $\Pr[X_i \geq x]$, while the latter happens with probability $q_{-i}(1)$. Since these two events are independent we have

$$\Pr[z_2 \geq x] = \frac{1}{2} \sum_{i \in \{1,2\}} q_{-i}(1) \Pr[X_i \geq x] \geq \frac{\theta(2)}{2} \sum_i \Pr[X_i \geq x] \quad .$$

where the last inequality follows from Proposition 6, although the proposition is trivial for $n = 2$. We can now continue analyzing $E[z]$ from before

$$\begin{aligned} E[z] &= \tau_2(1-\theta(2)) + (\tau_1 - \tau_2)(1-\theta(1)) + \int_{\tau_1}^{\infty} \Pr[z_1 \geq x]\, dx + \int_{\tau_2}^{\infty} \Pr[z_2 \geq x]\, dx \\ &\geq \tau_2(1-\theta(2)) + (\tau_1 - \tau_2)(1-\theta(1)) + \frac{\theta(1)}{2} \int_{\tau_1}^{\infty} \sum_i \Pr[X_i \geq x]\, dx + \frac{\theta(2)}{2} \int_{\tau_2}^{\infty} \sum_i \Pr[X_i \geq x]\, dx \quad . \end{aligned}$$

We note that although the $\theta(1)$ factor is not required in the third term of the last inequality, we include it so that the formulas can have the same formation as in the general formula of the next sections.

It remains to bound $\int_{\tau_k}^{\infty} \sum_i \Pr[X_i \geq x]$ for $k \in \{1, 2\}$. Recall that $OPT = E[\max_i X_i]$. Hence for every $k \in \{1, 2\}$ and since $\Pr[\max X_i \geq x] \leq 1$ we have

$$OPT = \int_0^{\tau_k} \Pr[\max X_i \geq x]\, dx + \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \leq \tau_k + \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \quad .$$

From $\tau_k = \alpha_k OPT$ and since $\Pr[\max X_i \geq x] \leq \sum_i \Pr[X_i \geq x]\, dx$ we then have

$$(1-\alpha_k)OPT \leq \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \leq \int_{\tau_k}^{\infty} \sum_i \Pr[X_i \geq x]\, dx \quad .$$

Therefore we get

$$\begin{aligned} E[z] &\geq \tau_2(1-\theta(2)) + (\tau_1 - \tau_2)(1-\theta(1)) + \frac{\theta(1)}{2} \int_{\tau_1}^{\infty} \sum_i \Pr[X_i \geq x]\, dx + \frac{\theta(2)}{2} \int_{\tau_2}^{\infty} \sum_i \Pr[X_i \geq x]\, dx \\ &\geq (\alpha_2 OPT)(1-\theta(2)) + (\alpha_1 - \alpha_2)OPT(1-\theta(1)) + \frac{\theta(1)}{2}(1-\alpha_1)OPT + \frac{\theta(2)}{2}(1-\alpha_2)OPT \\ &= OPT\left(\alpha_1 + \theta(1)(\frac{1+2\alpha_2-3\alpha_1}{2}) + \theta(2)(\frac{1-3\alpha_2}{2})\right) \quad . \end{aligned}$$

Therefore by choosing $\alpha_2 = 1/3$ and $\alpha_1 = 5/9$, the coefficients of $\theta(1)$ and $\theta(2)$ become zero, leading to the competitive ratio of $5/9 \simeq 0.55$. In the next section, we show how one can generalize the arguments to the case of n thresholds for arbitrary n.

6 $(1 - \frac{1}{e} \approx 0.63)$-Competitive Ratio Using n Thresholds

In this section we prove our main theorem. In particular, we invoke Algorithm Prophet Secretary with n distinct thresholds $\tau_1, \ldots, \tau_n$. The thresholds $\tau_1, \ldots, \tau_n$ that we consider are *non-adaptive* (i.e., Algorithm Prophet Secretary is oblivious to the history) and *non-increasing*. Intuitively, this is because as we move to the end of stream we should be more pessimistic and use lower thresholds to catch remaining higher values.

Formally, for every $k \in [n]$, we consider threshold $\tau_k = \alpha_k \cdot \mathsf{OPT}$ where the sequence $\alpha_1, \ldots, \alpha_n$ is non-increasing that is, $\alpha_1 \geq \alpha_1 \geq \ldots \geq \alpha_n$. We invoke Algorithm Prophet Secretary with these thresholds and analyze the competitive ratio of Algorithm Prophet Secretary with respect to coefficients α_k. Theorem 1 shows that there exists a sequence of coefficients α_k that leads to the competitive ratio of $(1 - 1/e) \approx 0.63$.

Proof of Theorem 1. We prove the theorem in two steps: First, we find a lower bound on $E[z]$ in terms of OPT and coefficients α_i. Second, we set coefficients α_k so that (1) α_1 becomes the competitive ratio of Algorithm Prophet Secretary and (2) α_1 converges to $1 - 1/e$, when n goes to infinity.

We start by proving the following auxiliary lemmas. In the first lemma, we find a lower bound for $\int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx$ based on $\mathsf{OPT} = E[\max_i X_i]$.

Lemma 8 $\int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \geq (1 - \alpha_k)\mathsf{OPT}$.

Proof : For an arbitrary $k \in [n]$ and since $\Pr[\max X_i \geq x] \leq 1$, we have

$$\begin{aligned}
\mathsf{OPT} = E\left[\max_i X_i\right] &= \int_0^{\infty} \Pr[\max X_i \geq x]\, dx \\
&= \int_0^{\tau_k} \Pr[\max X_i \geq x]\, dx + \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \\
&\leq \tau_k + \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \ .
\end{aligned}$$

Since, by definition $\tau_k = \alpha_k \cdot \mathsf{OPT}$, we have $(1 - \alpha_k) \cdot \mathsf{OPT} \leq \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx$. □

Next, to find a lower bound for $E[z]$, we first split it into two terms. Later, we find lower bounds for each one of these terms based on $\mathsf{OPT} = E[\max_i X_i]$.

Lemma 9 *Let $z = \sum_{k=1}^{n} z_k$ denote the value chosen by Algorithm* Prophet Secretary. *For z we have*

$$E[z] = \sum_{k=1}^{n} \int_0^{\tau_k} \Pr[z_k \geq x]\, dx + \sum_{k=1}^{n} \int_{\tau_k}^{\infty} \Pr[z_k \geq x]\, dx \ .$$

Proof : By Proposition 5 we have

$$\begin{aligned}
E[z] &= \int_0^{\infty} \Pr[z \geq x]\, dx = \sum_{k=1}^{n} \int_0^{\infty} \Pr[z_k \geq x]\, dx \\
&= \sum_{k=1}^{n} \int_0^{\tau_k} \Pr[z_k \geq x]\, dx + \sum_{k=1}^{n} \int_{\tau_k}^{\infty} \Pr[z_k \geq x]\, dx \ ,
\end{aligned}$$

where we use this fact that $z = \sum_{k=1}^{n} z_k$ because we only pick one number for a fixed sequence of drawn values and a fixed permutation and therefore, at most one of z_k's is non-zero. □

Lemma 10 $\sum_{k=1}^{n} \int_0^{\tau_k} \Pr[z_k \geq x]\, dx \geq \mathsf{OPT} \sum_{k=1}^{n} (1 - \theta(k))(\alpha_k - \alpha_{k+1})$.

Proof : Suppose $x \leq \tau_k$. Observe that the event $z_k \geq x$ occurs when Algorithm Prophet Secretary chooses a value at step k. In fact, since the thresholds are non-increasing, whatever we pick at the first k

steps would be at least x. Recall that for every $k \in [n]$, $\theta(k)$ is the probability that Algorithm Prophet Secretary does not choose a value from the first k steps. Hence, for every $k \in [n]$ and $x \leq \tau_k$ we have

$$\sum_{j \leq k} \Pr[z_j \geq x] = 1 - \theta(k) \ . \tag{4}$$

To simplify the notation, we assume that $\alpha_0 = \infty$ which means $\tau_0 = \infty$ and we let $\alpha_{n+1} = 0$ which means $\tau_{n+1} = 0$. Therefore we have

$$\sum_{k=1}^{n} \int_0^{\tau_k} \Pr[z_k \geq x] \, dx = \sum_{k=1}^{n} \int_{\tau_{n+1}}^{\tau_k} \Pr[z_k \geq x] \, dx \ .$$

Next, we use Equation (4) to prove the lemma as follows.

$$\begin{aligned}
\sum_{k=1}^{n} \int_0^{\tau_k} \Pr[z_k \geq x] \, dx &= \sum_{k=1}^{n} \int_{\tau_{n+1}}^{\tau_k} \Pr[z_k \geq x] \, dx \\
&= \sum_{k=1}^{n} \sum_{r=k}^{n} \int_{\tau_{r+1}}^{\tau_r} \Pr[z_k \geq x] \, dx \\
&= \sum_{r=1}^{n} \int_{\tau_{r+1}}^{\tau_r} \sum_{k=1}^{r} \Pr[z_k \geq x] \, dx \geq \sum_{r=1}^{n} \int_{\tau_{r+1}}^{\tau_r} (1 - \theta(r)) dx \\
&= \sum_{r=1}^{n} (1 - \theta(r))(\tau_r - \tau_{r+1}) = \mathrm{OPT} \cdot \sum_{k=1}^{n} (1 - \theta(k))(\alpha_k - \alpha_{k+1}) \ .
\end{aligned}$$

□

Lemma 11 $\sum_{k=1}^{n} \int_{\tau_k}^{\infty} \Pr[z_k \geq x] \, dx \geq \mathrm{OPT} \sum_k \frac{\theta(k)}{n} (1 - \alpha_k)$.

Proof : Recall that for every distribution D_i we draw a number X_i. Later, we randomly permute the numbers $X_1, \cdots, X_n$. Let the sequence of indices after the random permutation be $\pi_1, \ldots, \pi_n$ that is, at step k, number X_{π_k} for $\pi_k \in [n]$ is revealed.

Suppose $x \geq \tau_k$. We break the event $z_k > 0$ to n different scenarios depending on which index of the distributions $D_1, \cdots, D_n$ is mapped to index π_k in the random permutation. Let us consider the scenario in which Algorithm Prophet Secretary chooses the value drawn from a distribution i at step k. Such a scenario happens if (i) Algorithm Prophet Secretary does not choose a value from the first $k-1$ steps which are not drawn from i, and (ii) $X_i \geq \tau_k$. Observe that the two events are independent. Therefore, we have $\Pr[z_k \geq x] = \sum_i \Pr[\pi_k = i] \cdot \Pr[X_i \geq x] \cdot q_{-i}(k-1)$, where $q_{-i}(k)$ for every $i \in n$ and $k \in [n-1]$ is the probability that the following two events concurrently happen: (i) Algorithm Prophet Secretary does not choose a value from the first k elements, and (ii) none of the first k values are drawn from D_i. Since π_k is an index in the random permutation we obtain

$$\Pr[z_k \geq x] = \sum_i \Pr[\pi_k = i] \Pr[X_i \geq x] \cdot q_{-i}(k-1) = \frac{1}{n} \sum_i \Pr[X_i \geq x] \, q_{-i}(k-1) \ .$$

Using Proposition 6 and an application of the union bound we then have

$$\begin{aligned}\Pr[z_k \geq x] &= \sum_i \Pr[\pi_k = i] \cdot \Pr[X_i \geq x] \cdot q_{-i}(k-1) \\ &= \frac{1}{n} \sum_i \Pr[X_i \geq x] \cdot q_{-i}(k-1) \\ &\geq \frac{\theta(k)}{n} \cdot \sum_i \Pr[X_i \geq x] \geq \frac{\theta(k)}{n} \cdot \Pr\left[\max_i X_i \geq x\right] \ .\end{aligned}$$

Therefore, we obtain the following lower bound on $\sum_{k=1}^{n} \int_{\tau_k}^{\infty} \Pr[z_k \geq x]\, dx$.

$$\begin{aligned}\sum_{k=1}^{n} \int_{\tau_k}^{\infty} \Pr[z_k \geq x]\, dx &\geq \sum_k \int_{\tau_k}^{\infty} \frac{\theta(k)}{n} \Pr[\max X_i \geq x]\, dx \\ &= \sum_k \frac{\theta(k)}{n} \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \ .\end{aligned}$$

Finally, we use the lower bound of Lemma 8 for $\int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx$ to prove the lemma.

$$\begin{aligned}\sum_{k=1}^{n} \int_{\tau_k}^{\infty} \Pr[z_k \geq x]\, dx &\geq \sum_k \frac{\theta(k)}{n} \int_{\tau_k}^{\infty} \Pr[\max X_i \geq x]\, dx \\ &\geq \sum_k \frac{\theta(k)}{n} \cdot (1-\alpha_k) \cdot \mathsf{OPT} = \mathsf{OPT} \cdot \sum_k \frac{\theta(k)}{n} \cdot (1-\alpha_k) \ .\end{aligned}$$

□

Now we can plug in the lower bounds of Lemmas 10 and 11 into Lemma 9 to obtain a lower bound for $E[z]$.

Corollary 12 *Let $z = \sum_{k=1}^{n} z_k$ denote the value chosen by Algorithm* Prophet Secretary. *For z we have*

$$E[z] \geq \mathsf{OPT} \cdot (\alpha_1 + \sum_{k=1}^{n} \theta(k)(\frac{1}{n} - \frac{\alpha_k}{n} - \alpha_k + \alpha_{k+1})) \ .$$

We finish the proof of the theorem by proving the following claim.

Lemma 13 *The competitive ratio of Algorithm* Prophet Secretary *is at least α_1 which quickly converges to $1 - 1/e \approx 0.63$ when n goes to infinity.*

Proof : Using Corollary 12, for z we have

$$E[z] \geq \mathsf{OPT}\left(\alpha_1 + \sum_{k=1}^{n} \theta(k)\left(\frac{1}{n} - \frac{\alpha_k}{n} - \alpha_k + \alpha_{k+1}\right)\right) \ ,$$

which means that the competitive ratio depends on the probabilities $\theta(k)$'s. However, we can easily get rid of the probabilities $\theta(k)$'s by choosing α_k's such that for every k, $\left(\frac{1}{n} - \frac{\alpha_k}{n} - \alpha_k + \alpha_{k+1}\right) = 0$.

More formally, by starting from $\alpha_{n+1} = 0$ and choosing $\alpha_k = \frac{1+n\alpha_{k+1}}{1+n}$ for $k \leq n$, the competitive ratio of the algorithm would be α_1. Below, we show that when n goes to infinity, α_1 quickly goes to $1 - 1/e$ which means that the competitive ratio of Algorithm Prophet Secretary converges to $1 - 1/e \approx 0.63$.

First, we show by the induction that $\alpha_k = \sum_{i=0}^{n-k} \frac{n^i}{(1+n)^{i+1}}$. For the base case we have

$$\alpha_n = \frac{1 + n\alpha_{n+1}}{1+n} = \frac{1 + n \times 0}{1+n} = \frac{n^0}{(1+n)^1} \ .$$

Given $\alpha_{k+1} = \sum_{i=0}^{n-(k+1)} \frac{n^i}{(1+n)^{i+1}}$ we show the equality for α_k as follows.

$$\begin{aligned}
\alpha_k = \frac{1 + n\alpha_{k+1}}{1+n} &= \frac{1 + n(\sum_{i=0}^{n-(k+1)} \frac{n^i}{(1+n)^{i+1}})}{1+n} \\
&= \frac{n^0}{(1+n)^1} + \sum_{i=0}^{n-(k+1)} \frac{n^{i+1}}{(1+n)^{i+2}} \\
&= \sum_{i=0}^{n-k} \frac{n^i}{(1+n)^{i+1}} \ .
\end{aligned}$$

Now we are ready to show $\alpha_1 \geq 1 - 1/e$ when n goes to infinity.

$$\begin{aligned}
\lim_{n\to\infty} \alpha_1 = \lim_{n\to\infty} \sum_{i=0}^{n-1} \frac{n^i}{(n+1)^{i+1}} &= \lim_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{n-1} (1 - \frac{1}{n+1})^i \\
\approx \lim_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{n-1} e^{-i/n} &\approx \int_0^1 e^{-x} dx = 1 - 1/e \ .
\end{aligned}$$

□

7 Lower Bounds for Prophet Secretary and Minimization Variants of Classical Stopping Theory Problems

In this section we give our lower bounds for the prophet secretary problem and the minimization variants of the prophet inequality and prophet secretary problems. First, in Section 7.1 we show that no online algorithm for the prophet secretary problem can achieve a competitive ratio of 0.73. Later, in Section 7.2 we consider the minimization variant of the prophet inequality problem. We show that, even for the simple case in which numbers are drawn from *identical and independent distributions (i.i.d.)*, there is no constant competitive online algorithm for the minimization variants of the prophet inequality and prophet secretary problems.

7.1 0.73-Lower Bound of Prophet Secretary Problem

Here we first give a simple example which shows that no algorithm for the prophet secretary problem can achieve competitive ratio 0.75. Later, we improve this lower bound by giving an example for which no online algorithm for the prophet secretary problem can achieve a competitive ratio 0.73.

Lemma 14 *There is no online algorithm for the prophet secretary problem with competitive ratio* $0.75 + \epsilon$*, where* ϵ *is an arbitrary small positive number less than* 1.

Proof : It is known that no algorithm can guarantee a competitive ratio better than $0.5 + \epsilon$, where ϵ is an arbitrary small positive number less than 1. A hard example that shows this upper bound is as follow. We have two distributions. The first distribution always gives 1, and the second distribution gives $\frac{1}{\epsilon}$ with probability ϵ and 0 with probability $1 - \epsilon$. Observe that if we either accept the first number or reject it our expected outcome is 1. On the other hand, the offline optimum algorithm takes $\frac{1}{\epsilon}$ with probability ϵ and 1 with probability $1 - \epsilon$. Therefore, the expected outcome of the offline optimum algorithm is $\epsilon \cdot \frac{1}{\epsilon} + 1 \cdot (1 - \epsilon) = 2 - \epsilon$, and the competitive ratio is at most $\frac{1}{2-\epsilon} \leq 0.5 + \epsilon$.

The above example contains exactly two distributions. Thus, if the drawn numbers from these distributions arrive in a random order, with probability 0.5 the arrival order is the worst case order. This means that in the prophet secretary, with probability 0.5 the expected outcome of any algorithm on this example is at most 1, while the offline optimum algorithm is always $2 - \epsilon$. Therefore, there is no algorithm for the prophet secretary problem with competitive ratio better than $\frac{0.5\times 1 + 0.5\times(2-\epsilon)}{2-\epsilon} \leq 0.75 + \epsilon$.

□

Now we prove Theorem 3 that improves the above lower bound by showing that there is no algorithm for the prophet secretary problem with competitive ratio $\frac{11}{15} + \epsilon \approx 0.73 + \epsilon$.

Proof of Theorem 3. Suppose we have three distributions as follows. The first distribution always gives 1. The second distribution gives 2 with probability 0.5, and 2 with probability 0.5. The third distribution gives $\frac{1}{\epsilon}$ with probability ϵ and 0 with probability $1 - \epsilon$.

We observe that with probability ϵ the maximum is $\frac{1}{\epsilon}$, with probability $\frac{(1-\epsilon)}{2}$ the maximum is 2 and with probability $\frac{(1-\epsilon)}{2}$ the maximum is 1. Thus, the expected outcome of the offline optimum algorithm is

$$\epsilon \times \frac{1}{\epsilon} + \frac{(1-\epsilon)}{2} \times 2 + \frac{(1-\epsilon)}{2} \times 1 = 2.5 - 1.5\epsilon \ .$$

Next we analyze the expected outcome of an online algorithm for the secretary problem. We have 6 permutations for three distributions 1, 2, and 3. Figure 1 shows the expected outcome of an online algorithm for the secretary problem for each permutation. Permutation ijk means that the draw from the i-th distribution appears the first, the draw from the j-th distribution appears after that, and draw from the k-th distribution appears at the end. As we see in Figure 1, the expected outcome of any online algorithm for our example is at most $\frac{11+3.5\epsilon}{6}$. Thus, the competitive ratio of any online algorithm for the prophet secretary problem for $0 \leq \epsilon < \frac{2.5}{9}$ is bounded by

$$\frac{11 + 3.5\epsilon}{6 \times (2.5 - 1.5\epsilon)} = \frac{11 + 3.5\epsilon}{15 - 9\epsilon} = \frac{11 - 9\epsilon}{15 - 9\epsilon} + \frac{12.5\epsilon}{15 - 9\epsilon} < \frac{11}{15} + \epsilon \ .$$

7.2 Lower Bounds for Minimization Variants of Classical Stopping Theory Problems

In this section, we consider the minimization variant of the prophet secretary problem, in which we need to select one element from the input and we aim to minimize the value of the selected element. Indeed, we provide some hardness results for very simple and restricted cases of this problem. We extend our results to the minimization variant of the classical prophet inequality or minimization variant of secretary problem.

First, we consider the simple case that numbers are drawn from identical and independent distributions. In particular, we prove Theorem 4 that shows, even for the simple case of identical and independent distributions, there is no $\frac{(1.11)^n}{6}$ competitive online algorithm for the minimization variant of the prophet inequality problem. Since the input items come from identical distributions, independently, randomly reordering them

permutation	expected outcome (expanded form)	expected outcome (closed form)
123	$0.5 \times 2 + 0.5 \times \epsilon \times \frac{1}{\epsilon}$	1.5
132	$\epsilon \times \frac{1}{\epsilon} + (1-\epsilon) \times 0.5 \times 2$	$2-\epsilon$
213	$0.5 \times 2 + 0.5 \times 1$	1.5
231	$0.5 \times 2 + 0.5 \times \epsilon \times \frac{1}{\epsilon} + 0.5 \times (1-\epsilon) \times 1$	$2-\frac{\epsilon}{2}$
312	$\epsilon \times \frac{1}{\epsilon} + (1-\epsilon) \times 1$	$2-\epsilon$
321	$\epsilon \times \frac{1}{\epsilon} + (1-\epsilon) \times 0.5 \times 2$	$2-\epsilon$

Fig. 1: Outcomes of an online algorithm for the secretary problem.

do not change the distribution of the items in the input. Thus, this result holds for the minimization variant of the prophet secretary problem as well.

Later, we give more power to the online algorithm and let it to change its decision once; we call this model *online algorithm with one exchange*. In Theorem 15 we show that, for any large number C there is no C competitive online algorithm with one exchange for the minimization variant of the prophet inequality. In Corollary 16 and Corollary 17 we extend this hardness result to the minimization variant of prophet secretary problem and secretary problem, respectively.

Proof of Theorem 4. Suppose we have n identical distributions, each gives 0 with probability $\frac{1}{3}$, 1 with probability $\frac{1}{3}$ and 2^n with probability $\frac{1}{3}$. One can see that with probability $(\frac{1}{3})^n$, all of the numbers are 2^n and thus the minimum number is 2^n. Also, with probability $(\frac{2}{3})^n - (\frac{1}{3})^n$, there is no 0 and there is at least one 1, and thus, the minimum is 1. In all the other cases, the minimum is 0. Therefore, the expected outcome of the offline optimum algorithm is $(\frac{1}{3})^n \times 2^n + (\frac{2}{3})^n - (\frac{1}{3})^n < \frac{2^{n+1}}{3^n}$.

For this example, without loss of generality we can assume that any online algorithm accept 0 as soon as it appears, and also it does not accept 2^n except for the last item. Assume $i+1$ is the first time that the algorithm is willing to accept 1. The probability that we arrive in this point is $(\frac{2}{3})^i$ and the probability that we see 1 at that point is $\frac{1}{3}$. On the other hand, the probability that such an algorithm does not accept anything up to the last number and sees that the last number is 2^n is at least $(\frac{2}{3})^i(\frac{1}{3})^{n-i}$. Therefore, the expected outcome of the online algorithm is at least $(\frac{2}{3})^i\frac{1}{3} \times 1 + (\frac{2}{3})^i(\frac{1}{3})^{n-i} \times 2^n = \frac{2^i 3^{n-i-1}+3^i}{3^n}$. Thus, the competitive ratio is at least

$$\frac{2^i 3^{n-i-1} + 3^i}{2^{n+1}} = \frac{1}{6}(\frac{3^{n-i}}{2^{n-i}} + \frac{3^{i+1}}{2^n}) \ .$$

If $i \leq 0.73n$, the left term is at least $\frac{(1.11)^n}{6}$; otherwise, the right term is at least $\frac{(1.11)^n}{6}$.

Theorem 15 *For any large positive number C, there is no C-competitive algorithm for minimization prophet inequality with one exchange.*

Proof : Suppose we have three distributions as follows. The first distribution always gives 1. The second distribution gives $\frac{1}{\epsilon}$ with probability ϵ and gives $\frac{\epsilon}{1-\epsilon}$ with probability $1-\epsilon$. The third distribution gives $\frac{1}{\epsilon}$ with probability ϵ and gives 0 with probability $1-\epsilon$. We set ϵ later.

We observe that the minimum number is 0 with probability $1-\epsilon$, is $\frac{\epsilon}{1-\epsilon}$ with probability $\epsilon(1-\epsilon)$ and is 1 with probability ϵ^2. Thus, the expected outcome of the optimum algorithm is

$$(1-\epsilon) \times 0 + \epsilon(1-\epsilon) \times \frac{\epsilon}{1-\epsilon} + \epsilon^2 \times 1 = 2\epsilon^2 \ .$$

Now, we show that the outcome of any online algorithm with one exchange for this input is at least ϵ. In fact, this means that, the competitive ratio can not be less than $\frac{\epsilon}{2\epsilon^2} = \frac{1}{2\epsilon}$. If we set $\epsilon = \frac{1}{2C}$, this means that there is no C-competitive algorithm for the minimization prophet inequality with one exchange.

Recall that there is no uncertainty in the first number. If an algorithm withdraw the first number, it can take the outcome of either the second distribution or the third distribution. However, if we do this, with probability ϵ^2 the outcome is $\frac{1}{\epsilon}$. Thus, the expected outcome is at least ϵ as desired.

Now, we just need to show that if an algorithm does not select the first number, its expected outcome is at least ϵ. Observe that if this happens, then the algorithm has only the option of choosing either the second and the third distribution. We consider two cases. The first case is if the algorithm does not select the second number, then it must select the third number. Therefore, with probability ϵ, the outcome of the algorithm is 1 and thus, the expected outcome is at least ϵ. The second case occurs if the algorithm selects the second number and that is $\frac{\epsilon}{1-\epsilon}$. Therefore, with probability $1-\epsilon$, the outcome of the algorithm is $\frac{\epsilon}{1-\epsilon}$ and again the expected outcome is $\frac{\epsilon}{1-\epsilon} \times (1-\epsilon) = \epsilon$.

□

Finally, we show that even for the minimization variant of the prophet secretary there is no hope to get a constant competitive ratio.

Corollary 16 *For any large number C, there is no C-competitive algorithm for minimization prophet secretary with one exchange.*

Proof : In Theorem 15, we have three distributions. Thus, in the prophet secretary model the worst case order happen with probability $\frac{1}{6}$. Thus, the competitive ratio can not be more than $\frac{\frac{1}{6}\cdot\epsilon}{2\epsilon^2} = \frac{1}{12\epsilon}$. If we set $\epsilon = \frac{1}{12C}$, this essentially means that there is no C-competitive algorithm for the minimization prophet secretary with one exchange. □

Corollary 17 *For any large number C there is no C-competitive algorithm for the minimization secretary problem with one exchange.*

Proof : Suppose for the sake of contradiction that there exists an algorithm $\mathtt{Alg}$ for the minimization secretary problem which is C-competitive. Consider all possible realizations of the example in Theorem 15. Algorithm $\mathtt{Alg}$ is C-competitive when each of the these realizations comes in a random order. Therefore, Algorithm $\mathtt{Alg}$ is C-competitive when the input is a distribution over these realizations. This says that $\mathtt{Alg}$ is a C-competitive algorithm for the minimization prophet secretary, which contradicts Corollary 16 and completes the proof. □

Acknowledgments. We would like to thank Robert Kleinberg for fruitful discussion on early stages of this project.

References

[1] M. Ajtai, N. Megiddo, and O. Waarts. Improved algorithms and analysis for secretary problems and generalizations. *SIAM J. Discrete Math.*, 14(1):1–27, 2001.

[2] S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*. 2011.

[3] S. Alaei, M. Hajiaghayi, and V. Liaghat. Online prophet-inequality matching with applications to ad allocation. In *EC*, 2012.

[4] S. Alaei, M. Hajiaghayi, and V. Liaghat. The online stochastic generalized assignment problem. In *Approx*. Springer, 2013.

[5] S. Alaei, M. T. Hajiaghayi, V. Liaghat, D. Pei, and B. Saha. Adcell: Ad allocation in cellular networks. In *ESA*. 2011.

[6] D. Assaf, L. Goldstein, and E. Samuel-cahn. Ratio prophet inequalities when the mortal has several choices. *Ann. Appl. Prob*, 12:972–984, 2001.

[7] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. A knapsack secretary problem with applications. In *APPROX*, pages 16–28, 2007.

[8] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exch.*, 7(2):1–11, 2008.

[9] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA*, pages 434–443, 2007.

[10] M. Bateni, M. T. Hajiaghayi, and M. Zadimoghaddam. Submodular secretary problem and extensions. *ACM Transactions on Algorithms*, 9(4):32, 2013.

[11] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. 2010.

[12] E. B. Dynkin. The optimum choice of the instant for stopping a markov process. *Sov. Math. Dokl.*, 4:627–629, 1963.

[13] P. R. Freeman. The secretary problem and its extensions: a review. *Internat. Statist. Rev.*, 51(2):189–206, 1983.

[14] J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *J. Amer. Statist. Assoc.*, 61:35–73, 1966.

[15] K. S. Glasser, R. Holzsager, and A. Barron. The d choice secretary problem. *Comm. Statist. C—Sequential Anal.*, 2(3):177–199, 1983.

[16] M. T. Hajiaghayi, R. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *EC*, pages 71–80, 2004.

[17] M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, pages 58–65, 2007.

[18] N. Immorlica, R. D. Kleinberg, and M. Mahdian. Secretary problems with competing employers. In *WINE*, pages 389–400, 2006.

[19] D. P. Kennedy. Prophet-type inequalities for multi-choice optimal stopping. In *In Stoch. Proc. Applic.* 1978.

[20] R. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA*, pages 630–631, 2005.

[21] R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *STOC*, 2012.

[22] U. Krengel and L. Sucheston. Semiamarts and finite values. In *Bull. Am. Math. Soc.* 1977.

[23] U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. In *In Kuelbs, J., ed., Probability on Banach Spaces*. 1978.

[24] O. Lachish. $o(\log\log \mathrm{rank})$-competitive-ratio for the matroid secretary problem. In *FOCS*. 2014.

[25] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[26] R. J. Vanderbei. The optimal choice of a subset of a population. *Math. Oper. Res.*, 5(4):481–486, 1980.

[27] J. G. Wilson. Optimal choice and assignment of the best m of n randomly arriving items. *Stochastic Process. Appl.*, 39(2):325–343, 1991.

[28] Q. Yan. Mechanism design via correlation gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 710–719, 2011.