

Prophet Inequalities for Independent Random Variables from an Unknown Distribution

José R. Correa* Paul Dütting† Felix Fischer‡ Kevin Schewior§

Abstract

A central object in optimal stopping theory is the single-choice prophet inequality for independent, identically distributed random variables: Given a sequence of random variables X_1, \dots, X_n drawn independently from a distribution F , the goal is to choose a stopping time τ so as to maximize α such that for all distributions F we have $\mathbb{E}[X_\tau] \geq \alpha \cdot \mathbb{E}[\max_t X_t]$. What makes this problem challenging is that the decision whether $\tau = t$ may only depend on the values of the random variables X_1, \dots, X_t and on the distribution F . For quite some time the best known bound for the problem was $\alpha \geq 1 - 1/e \approx 0.632$ [Hill and Kertz, 1982]. Only recently this bound was improved by Abolhassani et al. [2017], and a tight bound of $\alpha \approx 0.745$ was obtained by Correa et al. [2017].

The case where F is unknown, such that the decision whether $\tau = t$ may depend only on the values of the first t random variables but not on F , is equally well motivated (e.g., [Azar et al., 2014]) but has received much less attention. A straightforward guarantee for this case of $\alpha \geq 1/e \approx 0.368$ can be derived from the solution to the secretary problem. Our main result is that this bound is tight. Motivated by this impossibility result we investigate the case where the stopping time may additionally depend on a limited number of samples from F . An extension of our main result shows that even with $o(n)$ samples $\alpha \leq 1/e$, so that the interesting case is the one with $\Omega(n)$ samples. Here we show that n samples allow for a significant improvement over the secretary problem, while $O(n^2)$ samples are equivalent to knowledge of the distribution: specifically, with n samples $\alpha \geq 1 - 1/e \approx 0.632$ and $\alpha \leq \ln(2) \approx 0.693$, and with $O(n^2)$ samples $\alpha \geq 0.745 - \varepsilon$ for any $\varepsilon > 0$.

*Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile. Email: correa@uchile.cl

†Department of Mathematics, London School of Economics, London, UK. Email: p.d.duetting@lse.ac.uk

‡School of Mathematical Sciences, Queen Mary University of London, London, UK. Email: felix.fischer@qmul.ac.uk

§Fakultät für Informatik, Technische Universität München, München, Germany; Département d'Informatique, École Normale Supérieure, PSL University, Paris, France. Email: kschewior@gmail.com. Supported by Conicyt Grant PII 20150140 and DAAD PRIME program.

1 Introduction

The theory of optimal stopping is concerned with what a computer scientist would call online algorithms, and the basic problem is one of sequential decision making with imperfect information about the future so as to maximize some reward or minimize some cost. Two canonical problems in the field are the *secretary problem* and the *prophet problem*. Both problems have over the past few years also received considerable attention from the theoretical computer science and algorithms community, particularly since they are closely related to the design of posted-price mechanisms in online sales.

In the *secretary problem* we are given n distinct, non-negative numbers from an unknown range. These numbers are presented in random order, and the goal is to stop at one of these numbers in order to maximize the probability with which we select the maximum. The problem has a surprisingly simple, and surprisingly positive, answer: by discarding a $1/e$ fraction of the numbers, and then selecting the first number that is greater than any of the discarded numbers, one is guaranteed to select the maximum with probability $1/e$ [e.g., Gilbert and Mosteller, 1966]. The guarantee of $1/e$ provided by this simple stopping rule is best possible, and remains best possible for example when numbers come from a uniform distribution with unknown and randomly chosen endpoints and are therefore correlated random variables [Berezovskiy and Gnedin, 1984; Ferguson, 1989]. When numbers are drawn independently from a single known distribution a better (and tight) guarantee of around 0.58 can be obtained [Gilbert and Mosteller, 1966].

In the *prophet problem* we are again shown n non-negative numbers, one at a time, but now these numbers are independent draws from known distributions and our goal is to maximize the expected value of the number on which we stop relative to the expected maximum value in hindsight. The two main results here concern the case where the distributions are distinct and the case where they are identical. For the former a tight bound of $1/2$ was given by Krengel and Sucheston [1977, 1978] and Samuel-Cahn [1984]. For the latter a lower bound of $1 - 1/e \approx 0.632$ due to Hill and Kertz [1982] was improved only very recently, first to 0.738 [Abolhassani et al., 2017] and then to 0.745 [Correa et al., 2017]. The bound of Correa et al. is in fact known to be tight due to a matching upper bound of Hill and Kertz [1982] and Kertz [1986].

An interesting variant of the prophet problem, for both identical and non-identical distributions, can be obtained by assuming that the distributions from which values are drawn are unknown. Despite being very natural (e.g., [Azar et al., 2014]), precious little is known about this variant.

Our Contribution We consider the prophet problem in which values are drawn independently from a single *unknown* distribution, and ask which approximation guarantees can be obtained relative to the expected maximum value in hindsight. This problem is interesting specifically for identical distributions, as here one could hope to learn something about later values from earlier ones. It seems challenging because, unlike in the case where the distribution is known and an optimal stopping rule can be obtained via backward induction, it is unclear what an optimal solution would look like.

A $1/e$ -approximation for our problem can be obtained in a relatively straightforward way by applying the secretary algorithm. The algorithm is guaranteed to stop on the maximum value with probability at least $1/e$, and one can show that this implies a $1/e$ -approximation relative to the expected maximum in hindsight. This analysis, however, seems crude and in particular does not take into account that we are rewarded also when we do not stop on the maximum value. Indeed, one would expect that the prophet objective is easier to achieve than the objective of the secretary problem.

Our main result shows that the straightforward guarantee of $1/e$ is in fact best possible in the prophet setting. The main difficulty in showing an impossibility result of this kind is that the set of possible stopping rules to which it applies is very rich. We will see, however, that for every stopping rule there exists a set $V \subseteq \mathbb{N}$ of arbitrary size and with an arbitrary gap between the largest and second-largest element on which the stopping rule is what we call value-oblivious: for random variables X_1, \dots, X_n with support V , the decision to stop at X_i when $X_i > \max\{X_1, \dots, X_{i-1}\}$ does not depend on the values of the random variables X_1, \dots, X_i but only on whether X_i is the largest among these values. We will then construct a distribution F with support V




Figure 1: Overview of results. The number of samples is displayed along the horizontal axis, the performance guarantee along the vertical axis. Lower bounds, shown as a solid line and two dots, result from stopping rules with a certain performance guarantee. Upper bounds, shown as dashed lines, correspond to impossibility results that no stopping rule can improve upon. The results for $o(n)$ and $O(n^2)$ samples are tight. With the exception of the upper bound of approximately 0.745, all results are new to this paper.

such that n values drawn independently from F are pairwise distinct with probability one and the expectation of their maximum is dominated by the largest value in V . The objective of the prophet problem on F is thus identical, up to a small error, to that of the secretary problem, and any stopping rule with a guarantee better than $1/e$ for the former would yield such a stopping rule for the latter. To understand why stopping rules must be value-oblivious it is useful to consider the special case where $n = 2$. In this case we may focus on stopping rules that always stop at X_2 whenever they have not stopped at X_1 , and every such stopping rule can be described by a function $p : \mathbb{R} \rightarrow [0, 1]$ such that $p(x)$ is the probability of stopping at X_1 when $X_1 = x$. By the Bolzano-Weierstrass theorem the infinite sequence $(p(n))_{n \in \mathbb{N}}$ contains a monotone subsequence and thus, for some $q \in [0, 1]$ and every $\varepsilon > 0$, a subsequence of values contained in the interval $[q - \varepsilon, q + \varepsilon]$. For random variables that only take values in the index set of that letter subsequence, the stopping rule will therefore stop at the first random random variable with what is essentially a fixed probability. When $n > 2$ the set of possible stopping rules becomes much richer, and identifying a set V on which a particular stopping rule is value-oblivious becomes much more challenging. Rather than the Bolzano-Weierstrass theorem, our proof uses the infinite version of Ramsey’s theorem [1930] to establish the existence of such a set.

Motivated by this impossibility result we then turn to the case where the stopping rule additionally has access to a limited number of samples from the distribution, which it may use in determining the stopping time. An extension of our main result shows that $o(n)$ samples are not enough to improve on the bound of $1/e$, the interesting case therefore is the one with $\Omega(n)$ samples. Our main result in this context is a simple but subtle stopping rule that achieves an approximation factor of $1 - 1/e \approx 0.632$ with $n - 1$ samples. We start by drawing $n - 1$ samples and using the maximum of these samples as a threshold for the first random variable. If the first random variable exceeds the threshold, we stop here. Otherwise we add the newly sampled value to the set of samples and then remove a random element from the resulting set. We then use the maximum of the new set as a threshold for the second random variable, and so on. While the procedure is easy to describe, its analysis is somewhat delicate. The key step is to show that the sets of random variables used to set the thresholds all behave like a set of $n - 1$ fresh samples. Thus the expected value collected from each random variable conditioned on accepting it equals the expected maximum value of n independent draws from the distribution, and the probability of accepting a random variable conditioned

on reaching it is exactly $1/n$. The approximation factor is then equal to the overall probability of stopping, which is at least $1 - 1/e$.

We complement the lower bound of $1 - 1/e$ with matching upper bounds for two different classes of algorithms that share specific properties of our algorithm. These bounds limit the types of approaches that could conceivably be used to go beyond a performance guarantee of $1 - 1/e$. We also provide a parametric upper bound for algorithms with access to γn samples for $\gamma \geq 0$. For algorithms that use at most n samples this upper bound is equal to $\ln(2) \approx 0.693$ and thus nearly tight.

Finally, we show how to get arbitrarily close to the optimal bound of 0.745 with $O(n^2)$ samples. The basic idea here is to mimic the optimal algorithm for known distributions, which uses a decreasing sequence of thresholds as determined by conditional acceptance probabilities, which are increasing over time. Our algorithm mirrors this approach using the corresponding quantiles of the empirical distribution function. It additionally skips a constant fraction of values at the beginning, and in our analysis we use the inequality of Dvoretzky, Kiefer, and Wolfowitz [1956] to show simultaneous concentration of all empirical quantiles. These two steps allow us to reduce the number of required samples from $O(n^4)$ to $O(n^2)$, relative to the obvious approach which uses all random variables and uses Chernoff and union bounds to show concentration. We provide evidence that any algorithm that achieves the optimal bound with $o(n^2)$ samples would have to use very different techniques.

In summary our results reveal a phase transition from secretary-like behavior to prophet-like behavior when going from $o(n)$ samples to $\Omega(n)$ samples, and show that $O(n^2)$ samples are equivalent to full knowledge of the distribution.

Further Related Work For early work on the classic single-choice prophet inequality in mathematics we refer the reader to a survey of Hill and Kertz [1992]. Starting from work of Hajiaghayi et al. [2007] prophet inequalities, and in particular extensions to richer feasibility domains, have seen a surge of interest in theoretical computer science (e.g., [Chawla et al., 2010; Alaei, 2014; Kleinberg and Weinberg, 2012; Feldman et al., 2015; Dütting and Kleinberg, 2015; Feldman et al., 2016; Rubinfeld, 2016; Rubinfeld and Singla, 2017; Dütting et al., 2017; Ehsani et al., 2018; Chawla et al., 2019]).

In theoretical computer science there is a relatively thin but important body of prior work on the case of unknown distributions. Most relevant for us is the aforementioned paper by Azar et al. [2014], which focuses on richer feasibility structures such as matching constraints and matroids, and earlier work by Babai et al. [2011], who consider a setting similar to ours but focus on a different objective, revenue maximization, apply different techniques, and obtain results that are qualitatively different from ours.

Related learning problems have also been studied in operations research and management science, but the types of problems, objectives, and techniques differ significantly from ours and typically involve regret minimization (see, e.g., the results of Goldenshluger and Zeevi [2017] and the recent survey of den Boer [2015]).

2 Preliminaries

Denote by \mathbb{N} the set of positive integers and let $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. For $i \in \mathbb{N}$, let $[i] = \{1, \dots, i\}$ and denote by \mathcal{S}_i the set of permutations of $[i]$.

Let $k \in \mathbb{N}_0$ and $n \in \mathbb{N}$. We consider (k, n) -stopping rules that sequentially observe random variables X_1, \dots, X_n and have access to samples S_1, \dots, S_k , and for each $i = 1, \dots, n$ decide whether to stop on X_i based on the values of X_1, \dots, X_i and S_1, \dots, S_k . We assume that X_1, \dots, X_n and S_1, \dots, S_k are independent and identically distributed, and respectively denote by f and F the probability density function and cumulative distribution function of their distribution. Formally, a (k, n) -stopping rule \mathbf{r} is a family of functions r_1, \dots, r_n where $r_i : \mathbb{R}_+^{k+i} \rightarrow [0, 1]$ for all $i = 1, \dots, n$. Here, $r_i(s_1, \dots, s_k, x_1, \dots, x_i)$ for $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$ is the probability of stopping at X_i conditioned on having received $S_1 = s_1, \dots, S_k = s_k$ as samples and $X_1 = x_1, \dots, X_i = x_i$ as values and not having stopped on any of X_1, \dots, X_{i-1} . The *stopping time* τ of a (k, n) -stopping rule \mathbf{r} ,

given S_1, \dots, S_k and X_1, \dots, X_n , is thus the random variable with support $\{1, \dots, n\} \cup \{\infty\}$ such that

$$\Pr[\tau = i \mid S_1 = s_1, \dots, S_k = s_k, X_1 = x_1, \dots, X_n = x_n] = \left(\prod_{j=1}^{i-1} (1 - r_j(s_1, \dots, s_k, x_1, \dots, x_j)) \right) \cdot r_i(s_1, \dots, s_k, x_1, \dots, x_i)$$

for all $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$.

For a given stopping rule we will be interested in the expected value $\mathbb{E}[X_\tau]$ of the variable at which it stops, where we use the convention that $X_\infty = 0$, and will measure its performance relative to the expected maximum $\mathbb{E}[\max\{X_1, \dots, X_n\}]$ of the random variables X_1, \dots, X_n . We will say that a stopping rule achieves approximation guarantee α , for $\alpha \leq 1$, if for any distribution, $\mathbb{E}[X_\tau] \geq \alpha \mathbb{E}[\max\{X_1, \dots, X_n\}]$.

For ease of exposition we will assume continuity of F in proving lower bounds and use discrete distributions to prove upper bounds. All results can be shown to hold in general by standard arguments, to break ties among random variables and to approximate a discrete distribution by a continuous one.

3 Sublinear Number of Samples

In this section we show that for $o(n)$ samples, the prophet problem with an unknown distribution behaves like the secretary problem. As we will see in Section 3.1, a straightforward baseline can be obtained from the optimal solution to the secretary problem, which discards a $1/e$ fraction of the values and then accepts the first value that exceeds the maximum of the discarded values. The algorithm does not require any samples, is guaranteed to stop at the maximum of the sequence with probability $1/e$, and can be shown to also provide a $1/e$ approximation for our objective. This analysis seems crude and in particular does not account for the fact that the prophet inequality is rewarded even when it does not stop on the maximum value of the sequence. Indeed the objective of the prophet problem seems easier to achieve than that of the secretary problem, and one would expect to be able to improve on the bound of $1/e$. Our main result, which we prove in Section 3.2, shows that this is not the case: the bound of $1/e$ is in fact best possible. This result continues to hold with $o(n)$ samples.

3.1 A $1/e$ -Approximation Without Samples

The following result translates the guarantee of $1/e$ for the secretary problem to a prophet inequality for independent random variables from an unknown distribution.

Theorem 3.1. *Let X_1, X_2, \dots, X_n be i.i.d. random variables drawn from an unknown distribution F . Then there exists a $(0, n)$ -stopping rule with stopping time τ such that*

$$\mathbb{E}[X_\tau] \geq \frac{1}{e} \cdot \mathbb{E}[\max\{X_1, X_2, \dots, X_n\}].$$

The result can be shown in a straightforward way, based on the idea that the realizations of the random variables X_1, \dots, X_n can be obtained by drawing n values from their common distribution and then permuting them uniformly at random. The classic analysis of the secretary problem [Ferguson, 1989] implies that *for each realization of the n draws*, the optimal stopping rule for this problem obtains the maximum value with probability $1/e$. It thus also obtains at least a $1/e$ fraction of the expected value of this maximum. We formalize this idea and prove Theorem 3.1 in the appendix.

3.2 A Matching Upper Bound

We proceed to show our main result: perhaps surprisingly, it is impossible to improve on the straightforward lower bound of $1/e$.

Theorem 3.2. *Let $\delta > 0$. Then there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$ and any $(0, n)$ -stopping rule with associated stopping time τ there exists a distribution F , not known to the stopping rule, such that when X_1, \dots, X_n are i.i.d. random variables drawn from F ,*

$$\mathbb{E}[X_\tau] \leq \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

The main difficulty in showing an impossibility result of this kind is that it applies to the set of all possible $(0, n)$ -stopping rules, which a priori is very rich. Indeed, recall that a $(0, n)$ -stopping rule \mathbf{r} is any family of functions r_1, \dots, r_n where $r_i : \mathbb{R}_+^i \rightarrow [0, 1]$ for all $i = 1, \dots, n$. Our main structural insight will be that we can restrict attention to stopping rules \mathbf{r} for which we can find arbitrarily large sets $V \subseteq \mathbb{N}$, such that for random variables X_1, \dots, X_n with support V , under the condition that X_1, \dots, X_i are pairwise distinct and $X_i > \max\{X_1, \dots, X_{i-1}\}$, and up to an arbitrarily small error ε , the probability of \mathbf{r} to stop on X_i does not depend on the *values* of any of the random variables X_1, \dots, X_i . This is made precise by the following definition.

Definition 1. *Let $\varepsilon > 0$ and $V \subseteq \mathbb{N}$. A stopping rule \mathbf{r} is called ε -value-oblivious on V if, for all $i \in [n]$, there exists a $q_i \in [0, 1]$ such that, for all pairwise distinct $v_1, \dots, v_i \in V$ with $v_i > \max\{v_1, \dots, v_{i-1}\}$, it holds that $r_i(v_1, \dots, v_i) \in [q_i - \varepsilon, q_i + \varepsilon]$.*

While value-obliviousness significantly restricts the expressiveness of a stopping rule, this restriction turns out to be essentially without loss when it comes to the ability of the class of all stopping rules to achieve a certain guarantee across all possible distributions: for any stopping rule and any $\varepsilon > 0$, there exists a stopping rule with the same guarantee that is ε -value-oblivious for some infinite set $V \subseteq \mathbb{N}$. This is made precise by the following lemma, which we prove in Section 3.3.

Lemma 3.3 (Main structural lemma). *Let $\varepsilon > 0$. If there exists a stopping rule with guarantee α , then there exists a stopping rule \mathbf{r} with guarantee α and an infinite set $V \subseteq \mathbb{N}$ such that \mathbf{r} is ε -value-oblivious on V .*

With Lemma 3.3 at hand it is not difficult to prove Theorem 3.2. For any stopping rule and an appropriate value of ε , we identify a stopping rule \mathbf{r} with the same performance guarantee that is ε -value-oblivious on an infinite set $V \subseteq \mathbb{N}$. We then define a distribution F with finite support $S \subseteq V$ such that (i) there is a large gap between the largest and second-largest elements of S , (ii) n independent draws from F are pairwise distinct with probability close to 1, (iii) \mathbf{r} is ε -value-oblivious on S , and (iv) the performance guarantee of \mathbf{r} on the distribution is dominated by the probability of selecting the largest element of S . By (i) and (ii) the prophet problem for the unknown distribution F is then equivalent up to a small error to a secretary problem, and by (iii) and (iv) \mathbf{r} behaves on F essentially like a stopping rule for the secretary problem. A performance guarantee for \mathbf{r} of more than $1/e$ would thus contradict the optimality of this bound for the secretary problem.

Proof of Theorem 3.2. Consider a $(0, n)$ -stopping rule with performance guarantee $1/e + \delta$. Set $\varepsilon = 1/n^2$. By Lemma 3.3 there exists a stopping rule \mathbf{r} with performance guarantee $1/e + \delta$ and an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{r} is ε -value-oblivious. Denote by τ the stopping time of \mathbf{r} . Let $v_1, \dots, v_{n^3}, u \in V$ be pairwise distinct such that $u \geq n^3 \max\{v_1, \dots, v_{n^3}\}$. For each $i \in [n]$, let

$$X_i = \begin{cases} v_1 & \text{w.p. } \frac{1}{n^3} \cdot \left(1 - \frac{1}{n^2}\right) \\ \vdots & \\ v_{n^3} & \text{w.p. } \frac{1}{n^3} \cdot \left(1 - \frac{1}{n^2}\right) \\ u & \text{w.p. } \frac{1}{n^2} \end{cases}.$$

We proceed to bound $\mathbb{E}[\max\{X_1, \dots, X_n\}]$ from below and $\mathbb{E}[X_\tau]$ from above. For $i \in [n]$, let $X_{(i)}$ denote the i th order statistic of X_1, \dots, X_n , such that $X_{(n)} = \max\{X_1, \dots, X_n\}$. Then

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] \geq \Pr[X_{(n)} = u] \cdot u = \frac{1 - o(1)}{n} \cdot u.$$

On the other hand, using that $X_{(n)} = u$ with probability at most $1/n$, we have that

$$\begin{aligned} \mathbb{E}[X_\tau] &= \Pr[X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot \mathbb{E}[X_\tau \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \\ &\quad + \Pr[X_{(n)} = u \wedge X_{(n-1)} = u] \cdot \mathbb{E}[X_\tau \mid X_{(n)} = u \wedge X_{(n-1)} = u] \\ &\quad + \Pr[X_{(n)} \neq u] \cdot \mathbb{E}[X_\tau \mid X_{(n)} \neq u] \\ &\leq \frac{1}{n} \cdot \left(\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot u \right. \\ &\quad \left. + \Pr[X_\tau \neq X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot O(n^{-3}) \cdot u \right) \\ &\quad + O(n^{-2}) \cdot u + 1 \cdot O(n^{-3}) \cdot u \\ &\leq \frac{1 + o(1)}{n} \cdot \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot u \\ &\leq \frac{1 + o(1)}{n} \cdot \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] \cdot u. \end{aligned}$$

To complete the proof we argue that

$$\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] \leq 1/e + o(1).$$

To see this assume for contradiction that $\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] \geq 1/e + \delta$ for some $\delta > 0$. Since \mathbf{r} is ε -value oblivious on V , for each $i \in [n]$ and distinct $s_1, \dots, s_i \in V$ with $s_i > \max\{s_1, \dots, s_{i-1}\}$, there is a probability q_i such that $r_i(s_1, \dots, s_i) \in [q_i - \varepsilon, q_i + \varepsilon)$. We can thus define a new $(0, n)$ -stopping rule $\hat{\mathbf{r}}$ with associated stopping time $\hat{\tau}$ such that for all $s_1, \dots, s_i \in V$, $\hat{r}_i(s_1, \dots, s_i) = q_i$ if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $\hat{r}_i(s_1, \dots, s_i) = 0$ otherwise. Note that $\hat{\mathbf{r}}$ bases its decision to stop *only* on the relative ranks of the values seen so far.

For the sake of the analysis, we think of \mathbf{r} and $\hat{\mathbf{r}}$ as being coupled in the following way. Let c_1, \dots, c_n be n i.i.d. draws from the uniform distribution on $[0, 1]$. For any $i \in [n]$ and $s_1, \dots, s_i \in V$, conditioned on $\tau \geq i$, we have $\tau = i$ if and only if $r_i(s_1, \dots, s_i) > c_i$, and, conditioned on $\hat{\tau} \geq i$, we have $\hat{\tau} = i$ if and only if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $\hat{r}_i(s_1, \dots, s_i) = q_i > c_i$.

Now consider any fixed sequence $s_1, \dots, s_n \in V$ of distinct realizations of X_1, \dots, X_n , respectively. For any $i \in [n]$, define ξ_i to be the event that occurs if and only if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $c_i \in [\min\{r_i(s_1, \dots, s_i), \hat{r}_i(s_1, \dots, s_i)\}, \max\{r_i(s_1, \dots, s_i), \hat{r}_i(s_1, \dots, s_i)\}]$. Note that then $\Pr[\xi_i] \leq \varepsilon$. Further note that $X_\tau = X_{(n)}$ and $X_{\hat{\tau}} \neq X_{(n)}$ implies that ξ_i occurs for some $i \in [n]$. Hence, by the union bound,

$$\Pr[X_{\hat{\tau}} = X_{(n)} \mid X_1 = s_1, \dots, X_n = s_n] \geq \Pr[X_\tau = X_{(n)} \mid X_1 = s_1, \dots, X_n = s_n] - n\varepsilon. \quad (1)$$

Since this is true pointwise for all distinct $s_1, \dots, s_n \in V$,

$$\begin{aligned} &\Pr[X_{\hat{\tau}} = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] \\ &\geq \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] - n\varepsilon \\ &\geq 1/e + \delta' \end{aligned}$$

for some $\delta' > 0$, where we have used (1) for the first inequality and $\varepsilon = 1/n^2$ for the second one. However, under the condition that $X_{(n)} = u$ and X_1, \dots, X_n are pairwise distinct, the relative ranks of X_1, \dots, X_n are distributed uniformly at random. Thus $\hat{\mathbf{r}}$, which only relies on relative ranks, selects the maximum with probability $1/e + \delta'$, in contradiction to the well-known upper bound of $1/e$ for the secretary problem [Ferguson, 1989, Section 2]. \square

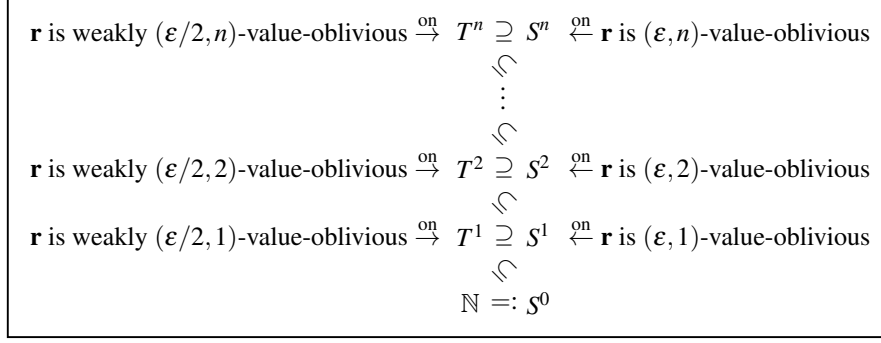


Figure 2: An overview of the infinite sets constructed in the proof of Lemma 3.3.

3.3 Proof of the Main Structural Lemma

We prove Lemma 3.3 through a sequence of steps that successively restrict the expressiveness of the stopping rules we have to consider. First we show a restriction to what we call order-oblivious rules, which in the decision to stop at random variable X_i , and conditioned of having reached X_i , may take into account the values of random variables X_1, \dots, X_{i-1} but not the order in which they were observed.

Definition 2. A stopping rule \mathbf{r} is order-oblivious if for all $j \in [n]$, all pairwise distinct $v_1, \dots, v_j \in \mathbb{R}_+$ and all permutations $\pi \in \mathcal{S}_{j-1}$, $r_i(v_1, \dots, v_j) = r_i(v_{\pi(1)}, \dots, v_{\pi(j-1)}, v_j)$.

The following result is very intuitive, but some care is required to prove it formally. We provide a proof in the appendix.

Lemma 3.4. If there exists a stopping rule with guarantee α , then there exists a stopping rule with guarantee α that is order-oblivious.

To further restrict the class of stopping rules from order-oblivious to value-oblivious ones we will now construct, for every order-oblivious rule \mathbf{r} and any $\varepsilon > 0$, an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{r} is ε -value-oblivious. The set V will depend on \mathbf{r} and will be obtained by starting from \mathbb{N} and identifying smaller and smaller subsets on which the behaviour of \mathbf{r} is more and more limited. By induction on $i \in [n]$ we will identify a set on which value-obliviousness holds with respect to the i th random variable. For each particular value of i , the latter will be shown in turn via a weaker form of obliviousness only to the value of X_i . We need the following definitions.

Definition 3. Consider a stopping rule \mathbf{r} . Let $\varepsilon > 0$, $i \in [n]$, and $V \subseteq \mathbb{N}$. Then \mathbf{r} is

- (a) (ε, i) -value-oblivious on V if there exists $q \in [0, 1]$ such that, for all pairwise distinct $v_1, \dots, v_i \in V$ with $v_i > \max\{v_1, \dots, v_{i-1}\}$, it holds that $r_i(v_1, \dots, v_i) \in [q - \varepsilon, q + \varepsilon]$;
- (b) weakly (ε, i) -value-oblivious on V if, for all pairwise distinct $v_1, \dots, v_{i-1} \in V$, there exists a $q \in [0, 1]$ such that, for all $v_i \in V$ with $v_i > \max\{v_1, \dots, v_{i-1}\}$, it holds that $r_i(v_1, \dots, v_i) \in [q - \varepsilon, q + \varepsilon]$.

Note that weak (ε, i) -value-obliviousness is indeed weaker than (ε, i) -value-obliviousness because the order of quantification has been reversed. Under the condition that $X_i > \max\{X_1, \dots, X_{i-1}\}$ the former requires the conditional probability of stopping at X_i to be independent of X_i , the latter also of X_1, \dots, X_{i-1} . Note further that (ε, i) -value-obliviousness for all $i \in [n]$ is equivalent to ε -value-obliviousness.

In establishing weak (ε, i) -value-obliviousness and (ε, i) -value-obliviousness for a particular value of i we will appeal to the infinite version of Ramsey's theorem to show the existence of an appropriate set V .

Lemma 3.5 (Ramsey [1930]). Let $c, d \in \mathbb{N}$, and let H be an infinite complete d -uniform hypergraph whose hyperedges are colored with c colors. Then there exists an infinite complete d -uniform sub-hypergraph of H that is monochromatic.

Proof of Lemma 3.3. Consider a stopping rule \mathbf{r} with guarantee α . By Lemma 3.4, it is without loss of generality to assume that \mathbf{r} is order-oblivious. We fix $\varepsilon > 0$ for the entire proof and show by induction on $j \in [n]$ that there exists an infinite set $S^j \subseteq \mathbb{N}$ such that, for all $i \in [j]$, \mathbf{r} is (ε, i) -value-oblivious on S^j . This suffices to show the claim, as for $j = n$ it implies that the stopping rule \mathbf{r} is (ε, j) -value-oblivious on S^n for all $j \in \mathbb{N}$, and hence ε -value-oblivious on S^n .

As $S^0 = \mathbb{N}$ satisfies the induction hypothesis for $j = 0$, we proceed to show it for $j = k > 0$ assuming that it is true for $j < k$. First observe that we only need to find an infinite set $S^k \subseteq S^{k-1}$ such that \mathbf{r} is (ε, k) -value-oblivious on S^k , because it follows from the induction hypothesis that S^k , as a subset of S^{k-1} , is (ε, i) -value-oblivious on S^i for all $i \in [k-1]$. We construct such a set in two steps. First we show that there exists an infinite set $T^k \subseteq S^{k-1}$ such that \mathbf{r} is weakly $(\varepsilon/2, k)$ -value-oblivious on T^k , formalized as Claim 1 below. Then we show that there exists an infinite set $S^k \subseteq T^k$ such that \mathbf{r} is (ε, k) -value-oblivious on S^k , formalized as Claim 2 below. See Figure 2 for a global overview.

Claim 1. There exists an infinite set $T^k \subseteq S^{k-1}$ such that \mathbf{r} is weakly $(\varepsilon/2, k)$ -value-oblivious on T^k .

We construct $T^k = \{t_1^k, t_2^k, \dots\}$ by induction. For the induction base to work, set $t_0^k := -1$. We define a set U_ℓ^k for any $\ell \geq 0$ such that $S^{k-1} \supseteq U_0^k \supseteq U_1^k \supseteq \dots$ and then set $t_\ell^k := \min U_\ell^k$ for all $\ell \geq 1$. Let $U_0^k := S^{k-1}$. For $\ell \geq 1$, U_ℓ^k is itself determined using another induction. Let $h = \binom{\ell-1}{k-1}$. We construct $h+1$ sets $U_{\ell,0}^k \supseteq U_{\ell,1}^k \supseteq \dots \supseteq U_{\ell,h}^k$, starting from $U_{\ell,0}^k := U_{\ell-1}^k \setminus \{t_{\ell-1}^k\}$, and then set $U_\ell^k := U_{\ell,h}^k$. Note that $h = 0$ when $\ell \leq k-1$, meaning that we have $U_\ell^k = U_{\ell,0}^k$ for these values of ℓ . For instance, assuming $\ell = 1 \leq k-1$ we have $U_1^k = U_{1,0}^k = U_0^k \setminus \{t_0^k\} = U_0^k = S^{k-1}$, so $t_1^k = \min U_1^k$ is the smallest value in S^{k-1} . Then, assuming $\ell = 2 \leq k-1$ we have $U_2^k = U_{2,0}^k = U_1^k \setminus \{t_1^k\}$, so $t_2^k = \min U_2^k$ is the second smallest value in S^{k-1} , and so on. In other words, $t_1^k, t_2^k, \dots, t_{k-1}^k$ are simply the $k-1$ smallest values in S^{k-1} .

When $\ell \geq k$, we have $h > 0$. Denote by P_1, \dots, P_h in arbitrary order the subsets of $\{t_1^k, t_2^k, \dots, t_{\ell-1}^k\}$ of cardinality $k-1$. To obtain $U_{\ell,i}^k$ for $i \in [h]$ consider $P_i = \{p_1, \dots, p_{k-1}\}$. Toward the first, and rather trivial, application of Lemma 3.5 consider an edge-colored complete 1-uniform hypergraph H_i with vertex set $U_{\ell,i-1}^k$, where for all $v \in U_{\ell,i-1}^k$, the singleton hyperedge containing v is colored with color $\lfloor r_k(p_1, \dots, p_{k-1}, v)/\varepsilon \rfloor$. Note that at most $\lfloor 1/\varepsilon \rfloor + 1$ colors are used in total. Now, by Lemma 3.5 with $c = \lfloor 1/\varepsilon \rfloor + 1$ and $d = 1$, there exists an infinite set of vertices that induces a monochromatic sub-hypergraph of H_i , and we define $U_{\ell,i}^k$ to be such a set of vertices. See Figure 3 for an overview of this construction.

It remains to be shown that \mathbf{r} is indeed weakly $(\varepsilon/2, k)$ -value-oblivious on T^k . Recall that this means that, for all distinct $v_1, \dots, v_{k-1} \in T^k$, there exists a $q \in [0, 1]$ such that for all $v_k \in T^k$ with $v_k > \max\{v_1, \dots, v_{k-1}\}$ we have $r_k(v_1, \dots, v_k) \in [q - \varepsilon/2, q + \varepsilon/2]$. So consider such distinct $v_1, \dots, v_{k-1} \in T^k$. Observe that there exists a smallest ℓ such that $\{v_1, \dots, v_{k-1}\} \subseteq \{t_1^k, t_2^k, \dots, t_\ell^k\}$. So for P_1, \dots, P_h being the sets of size $k-1$ considered for the construction of $U_{\ell+1}^k$, there must exist $i \in [h]$ such that $\{v_1, \dots, v_{k-1}\} = P_i$. Now for u being the color of the monochromatic infinite sub-hypergraph H'_i of the corresponding H_i selected, we set $q := u \cdot \varepsilon + \varepsilon/2$. By our choice of ℓ , for any $v_k \in T^k$ with $v_k > \max\{v_1, \dots, v_{k-1}\}$ we have $v_k = t_{\ell'}^k$ for some $\ell' > \ell$. So $v_k \in U_{\ell'}^k$, implying $v_k \in U_{\ell+1}^k \supseteq U_\ell^k$ and in particular $v_k \in U_{\ell+1,i}^k \supseteq U_{\ell+1}^k$. Therefore $v_k \in H'_i$. By construction of H_i , that means that $r_k(v_{\pi(1)}, \dots, v_{\pi(k-1)}, v_k) \in [q - \varepsilon/2, q + \varepsilon/2]$ for some permutation $\pi \in \mathcal{S}_{k-1}$. But since \mathbf{r} is order-oblivious, $r_k(v_1, \dots, v_k) = r_k(v_{\pi(1)}, \dots, v_{\pi(k-1)}, v_k)$. Thus we have shown that \mathbf{r} is weakly $(\varepsilon/2, k)$ -value-oblivious on T^k .

We note that the case $k = 1$ is degenerate in that the only set of size $k-1$ considered is the empty set, but the same construction and arguments work here. By construction $T^k \subseteq S^{k-1}$, so we have proved the claim.

Claim 2. There is an infinite set $S^k \subseteq T^k$ such that \mathbf{r} is (ε, k) -value-oblivious

Toward the second, and less trivial, application of Lemma 3.5, we construct a complete $(k-1)$ -uniform hypergraph H with vertex set T^k . Consider any set $\{v_1, \dots, v_{k-1}\} \subseteq T^k$ of cardinality $k-1$. By Claim 1, there exists $q \in [0, 1]$ such that, for all $v_k \in T^k$, we have $r_k(v_1, \dots, v_k) \in [q - \varepsilon/2, q + \varepsilon/2]$. Note that there

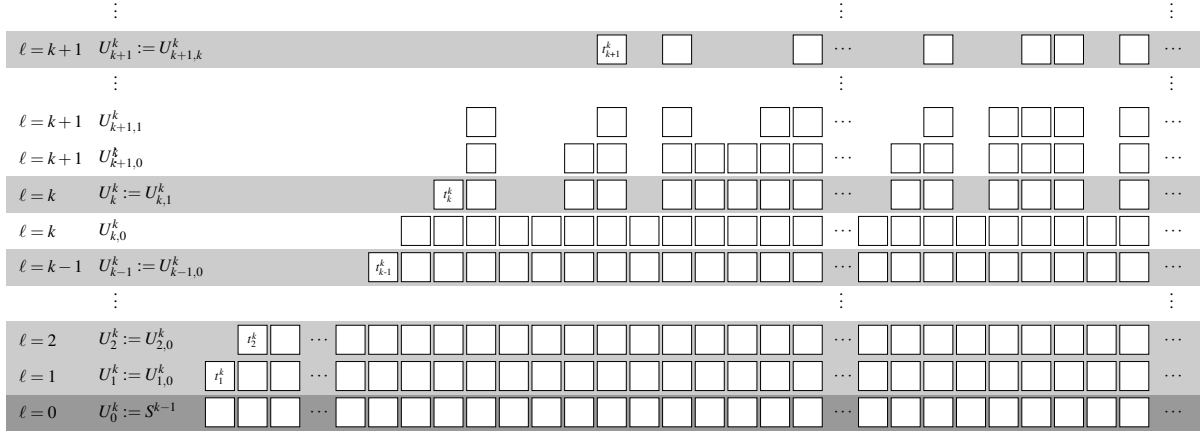


Figure 3: An overview of the sets constructed in the proof of Claim 1. Each of the elements in S^{k-1} corresponds to a column, which are arranged in increasing order. An element is included in a set (corresponding to a row) if there is a small square box in the corresponding cell.

exists a unique $u \in \mathbb{N}$ such that $u \cdot \varepsilon \in [q - \varepsilon/2, q + \varepsilon/2)$, and color the hyperedge $\{v_1, \dots, v_{k-1}\}$ of H with color u .

Note that for every color u , for every hyperedge $\{v_1, \dots, v_{k-1}\}$ with color u , and for every $v_k \in T^k$, $r_i(v_1, \dots, v_{k-1}, v_k) \in [u \cdot \varepsilon - \varepsilon, u \cdot \varepsilon + \varepsilon)$. For any set $V \subseteq T^k$ that induces a monochromatic sub-hypergraph of H we thus have that \mathbf{r} is (ε, k) -value-oblivious on V . By Lemma 3.5 with $c = \lfloor 1/\varepsilon \rfloor + 1$ and $d = k - 1$ there exists an infinite set that induces a monochromatic sub-hypergraph, and letting S^k be such a set completes the proof of the claim. \square

3.4 Extension of the Upper Bound to $o(n)$ Samples

We conclude this section by showing that even with $o(n)$ samples the guarantee of $1/e$ is still best possible.

Corollary 3.6. *Let $\delta > 0$ and $f : \mathbb{N} \rightarrow \mathbb{N}$ with $f(n) = o(n)$. Then there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$ and any $(f(n), n)$ -stopping rule with associated stopping time τ there exists a distribution F , not known to the stopping rule, such that when X_1, \dots, X_n are i.i.d. random variables drawn from F ,*

$$\mathbb{E}[X_\tau] \leq \left(\frac{1}{e} + \delta \right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

We give some intuition for why this is true: Assume there exists a $(o(n), n)$ -stopping rule \mathbf{r} with guarantee bounded away from $1/e$. Then we could obtain a $(0, n)$ -stopping rule \mathbf{r}' by interpreting (for suitable n') the first $o(n')$ values as samples, the following n' values as actual values on which the rule may stop, and then running \mathbf{r} in this setting. As we can choose $n' = (1 - o(1)) \cdot n$, the expected maximum of n and n' draws (from any distribution) are identical up to a $(1 - o(1))$ factor, so the guarantee of \mathbf{r} carries over to \mathbf{r}' , contradicting Theorem 3.2. We give a short formal proof based on Theorem 4.5 in Appendix C.

4 Linear Number of Samples

The previous section has revealed a strong impossibility: even with $o(n)$ samples it is impossible to improve over the straightforward lower bound of $1/e \approx 0.368$ achieved by the well-known optimal stopping rule for the secretary problem. We proceed to show that there is a sharp phase transition when going from $o(n)$ samples to $\Omega(n)$ samples, by giving an algorithm that uses as few as $n - 1$ samples and improves the lower bound from $1/e$ to $1 - 1/e \approx 0.632$. We also show that the bound of $1 - 1/e$ is in fact tight for two different classes of algorithms that share certain features of our algorithm. This illustrates that our analysis is tight

and limits the types of approaches that could conceivably be used to go beyond $1 - 1/e$. We also show a parametric upper bound for algorithms that use γn samples for any $\gamma \geq 0$. For algorithms that use at most n samples this bound is equal to $\ln(2) \approx 0.693$ and thus nearly tight.

4.1 Warm-Up: A $1/2$ -Approximation with $n - 1$ Samples

To gain some intuition let us first consider the natural approach to sample $n - 1$ values S_1, \dots, S_{n-1} from F and to use the maximum of these samples as a uniform threshold for all of the random variables X_1, \dots, X_n , accepting the first random variable that exceeds the threshold. It is not difficult to see that the expected value we collect from any random variable X_i conditioned on stopping at that random variable is at least $\mathbb{E}[\max\{X_1, \dots, X_n\}]$, since under this condition X_i is the maximum of at least n i.i.d. random variables. We can thus understand the approximation guarantee provided by this approach by understanding the probability that it stops on some random variable. It turns out that this probability, and hence the approximation guarantee, is $1/2 + 1/(4n - 2)$. A more detailed analysis, which we provide in Appendix D, also reveals that an improvement over the bound of roughly $1/2$ is impossible with a uniform threshold, even if this threshold is chosen with knowledge of the distribution F .

4.2 A $(1 - 1/e)$ -Approximation with $n - 1$ Samples

We proceed to show that it is indeed possible to obtain an improved bound of $1 - (1 - 1/n)^n \geq 1 - 1/e \approx 0.632$ with just $n - 1$ samples. Our algorithm improves over the naïve approach that obtains a factor $1/2$ by increasing the probability that we stop at all, while maintaining the property that the expected value that we collect when we do stop is at least $\mathbb{E}[\max\{X_1, \dots, X_n\}]$.

Theorem 4.1. *Let X_1, X_2, \dots, X_n be i.i.d. random variables from an unknown distribution F . Then there exists an $(n - 1, n)$ -stopping-rule with stopping time τ such that*

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

Let us first convince ourselves that the bound would be straightforward to achieve if we were given access to $n(n - 1) \in \Theta(n^2)$ samples. In this case an approximation to the bound could be obtained from a result of Ehsani et al. [2018] by observing that $\Theta(n^2)$ samples provide a good approximation of the distribution of the maximum and by setting as a threshold the $1/e$ -quantile of that distribution. Here we take a different route that yields the bound exactly and that, more importantly, can be developed further to work with only $n - 1$ samples. To this end we partition the $n(n - 1)$ samples into n sets of size $n - 1$ each, and use the maximum of the i th set as a threshold for the i th random variable. Upon acceptance of any random variable, that random variable would have a value equal to the expected maximum of n i.i.d. random variables, which is equal to $\mathbb{E}[\max\{X_1, \dots, X_n\}]$. Conditioned on reaching the i th random variable it would be accepted with probability $1/n$, for an overall probability of acceptance of $\sum_{i=1}^n (1 - 1/n)^{i-1} \cdot 1/n = 1 - (1 - 1/n)^n$.

Algorithm 1 mimics this approach, but instead of using $n - 1$ fresh samples for each of the n random variables constructs $n - 1$ fresh-looking samples for each of the n random variables from a *single* set $\{S_1, \dots, S_{n-1}\}$ of $n - 1$ samples. For the first random variable X_1 the algorithm uses a threshold equal to the maximum of the $n - 1$ samples. If $X_1 \geq \max\{S_1, \dots, S_{n-1}\}$, the algorithm stops. Otherwise it adds X_1 to the set of samples, picks one of the elements in $\{S_1, \dots, S_{n-1}, X_1\}$ uniformly at random, and drops this element from the set. The algorithm then continues in the same way, by using the maximum of the set thus obtained as a threshold for the next random variable, and updating the set when a random variable fails to exceed its threshold.

To analyze the algorithm it will be useful to consider a sequence j_1, \dots, j_{n-1} of random variables drawn independently and uniformly from $[n]$. Then, for $i = 1, \dots, n$, define variables R_1^i, \dots, R_n^i recursively as

Algorithm 1: Fresh looking samples

Data: Sequence of i.i.d. random variables X_1, \dots, X_n sampled from an unknown distribution F ,
sample access to F

Result: Stopping time τ

$\tau \leftarrow n + 1$

$S_1, \dots, S_{n-1} \leftarrow n - 1$ independent samples from F

$S \leftarrow \{S_1, \dots, S_{n-1}\}$

for $t = 1, \dots, n$ **do**

if $X_t \geq \max\{S\}$ **then**

$\tau \leftarrow t$

 Break

else

$S \leftarrow S \cup \{X_t\}$

$Z \leftarrow$ value from S , chosen uniformly at random

$S \leftarrow S \setminus \{Z\}$

return τ

follows

$$R_\ell^1 = \begin{cases} S_\ell & \text{for } \ell = 1, \dots, n-1 \\ X_1 & \text{for } \ell = n \end{cases} \quad \text{and} \quad R_\ell^i = \begin{cases} R_\ell^{i-1} & \text{for } \ell = 1, \dots, j_{i-1} - 1 \\ R_{\ell+1}^{i-1} & \text{for } \ell = j_{i-1}, \dots, n-1 \\ X_i & \text{for } \ell = n \end{cases} \quad \text{for } t > 1.$$

Then the random variable $X_i = R_n^i$ and the threshold that we set for this random variable is $\max\{R_1^i, \dots, R_{n-1}^i\}$. Denote by ξ_i the event that random variable X_i exceeds the threshold that we set for it, that is $R_n^i = \max\{R_1^i, \dots, R_n^i\}$. An important observation is that ξ_i is independent from ξ_1, \dots, ξ_{i-1} .

Lemma 4.2. For every $i \in [n]$, $\Pr[\xi_i \cap (\bigcap_{j < i} \neg \xi_j)] = \Pr[\xi_i] \cdot \prod_{j < i} \Pr[\neg \xi_j]$.

Proof. It suffices to show that for all $t = 1, \dots, i-1$, the event $\xi_i \cap (\bigcap_{j=t+1}^{i-1} \neg \xi_j)$ is independent of the event $\neg \xi_t$, i.e.,

$$\Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid \neg \xi_t \right] = \Pr \left[\xi_i \cap \bigcap_{j=t+1}^{i-1} \neg \xi_j \right].$$

We claim that

$$\begin{aligned} \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \right] &= \sum_{\ell=1}^n \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_\ell^t = \max\{R_1^t, \dots, R_n^t\} \right] \cdot \Pr \left[R_\ell^t = \max\{R_1^t, \dots, R_n^t\} \right] \\ &= \frac{n-1}{n} \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_n^t < \max\{R_1^t, \dots, R_n^t\} \right] \\ &\quad + \frac{1}{n} \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_n^t = \max\{R_1^t, \dots, R_n^t\} \right] \\ &= \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_n^t < \max\{R_1^t, \dots, R_n^t\} \right] \end{aligned}$$

$$= \Pr \left[\xi_i \cap \left(\bigcap_{j=i+1}^{i-1} \neg \xi_j \right) \mid \neg \xi_i \right].$$

Indeed, the first equality can be obtained by distinguishing the index where the maximum is attained. For the second equality observe that the first probability on its left-hand side is the same for all values of ℓ because the events are independent of the choice of ℓ and conditioning is symmetric, and that the maximum is attained with probability $1/n$ at each index. This establishes the claim. \square

Proof of Theorem 4.1. The value $\mathbb{E}[X_\tau]$ obtained by Algorithm 1 can be written by summing over all possible stopping times $i = 1, \dots, n$ the product of the probability of stopping at $X_i = R_n^i$ and the expectation of X_i upon stopping, i.e.,

$$\mathbb{E}[X_\tau] = \sum_{i=1}^n \Pr[\xi_i \wedge \neg \xi_{i-1} \wedge \dots \wedge \neg \xi_1] \cdot \mathbb{E}[R_n^i \mid \xi_i \wedge \neg \xi_{i-1} \wedge \dots \wedge \neg \xi_1].$$

By Lemma 4.2, and since for each $i \in [n]$, the set $\{R_1^i, \dots, R_n^i\}$ is a set of n i.i.d. random variables,

$$\Pr[\xi_i \wedge \neg \xi_{i-1} \wedge \dots \wedge \neg \xi_1] = \left(1 - \frac{1}{n}\right)^{i-1} \frac{1}{n}.$$

Since R_n^i is independent of ξ_1, \dots, ξ_{i-1} , and using again that $\{R_1^i, \dots, R_n^i\}$ is a set of n i.i.d. random variables,

$$\mathbb{E}[R_n^i \mid \xi_i \wedge \neg \xi_{i-1} \wedge \dots \wedge \neg \xi_1] = \mathbb{E}[R_n^i \mid \xi_i] = \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

Thus

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

as claimed. \square

It is worth noting that the approach used by Algorithm 1 to create fresh-looking sets of samples is not unique. To determine a threshold for random variable X_i we may for example, independently at each step i , select uniformly at random a subset of size $n - 1$ from the set $\{S_1, \dots, S_{n-1}, X_1, \dots, X_{i-1}\}$, and use the maximum value in that set as a threshold. An independence result analogous to Lemma 4.2 could then be used to argue that the sequence of thresholds again provides a performance guarantee of $1 - 1/e$. An appealing property of the particular approach used by Algorithm 1 is that the sequence of thresholds it creates is non-increasing.

4.3 Going Beyond $1 - 1/e$

We proceed to show an upper bound of $1 - 1/e$ for two different classes of algorithms that share certain features of our algorithm. This shows that our analysis of Algorithm 1 is tight and limits the class of algorithms that could conceivably go beyond $1 - 1/e$.

The first upper bound applies to algorithms for which the probability of stopping at the i th random variable conditioned on reaching it is independent of i . This is true for Algorithm 1 since, by Lemma 4.2, $\Pr[\xi_i \mid (\bigcap_{j<i} \neg \xi_j)] = \Pr[\xi_i] = 1/n$. The upper bound applies even in the case where the distribution F is known, and to stopping rules that like Algorithm 1 use dependent thresholds. We provide a proof of this result in the appendix.

Proposition 4.3. *Let $\varepsilon > 0$. Then there exists $n \in \mathbb{N}$ and a distribution F such that for any stopping time τ for which $\Pr[\tau = i \mid \tau > i - 1]$ is independent of i ,*

$$\mathbb{E}[X_\tau] \leq \left(1 - \frac{1}{e} + \varepsilon\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

The second upper bound applies to any algorithm that like Algorithm 1 has access to $n - 1$ samples S_1, \dots, S_{n-1} from the underlying distribution and satisfies the following two natural conditions: (i) if the value of the first random variable X_1 is greater than all $n - 1$ samples, the algorithm stops; and (ii) conditioned on reaching X_i , the probability of stopping at X_i is nondecreasing in i . The proof of this result can be found in the appendix.

Proposition 4.4. *Let $\varepsilon > 0$. Then there exists $n \in \mathbb{N}$ and a distribution F such that for any $(n - 1, n)$ -stopping rule with stopping time τ that satisfies conditions (i) and (ii),*

$$\mathbb{E}[X_\tau] \leq \left(1 - \frac{1}{e} + \varepsilon\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

4.4 Close to Tight Upper Bound

While an improvement over the bound of $1 - 1/e \approx 0.632$ remains possible via more complicated stopping rules, such an improvement cannot go beyond $\ln(2) \approx 0.693$. This is a consequence of the following strengthening of Theorem 3.2, which provides a parametric upper bound for stopping rules that have access to γn samples for some $\gamma \geq 0$ and is proven in the appendix.

Theorem 4.5. *Let $\delta > 0$, $\gamma \in \mathbb{Q}$. Then there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$ and any $(\gamma n, n)$ -stopping rule with associated stopping time τ there exists a distribution F , not known to the stopping rule, with the following property. When X_1, \dots, X_n are i.i.d. random variables drawn from F , we have*

$$\mathbb{E}[X_\tau] \leq (b(\gamma) + \delta) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where

$$b(\gamma) = \begin{cases} \frac{1+\gamma}{e} & \text{if } \frac{1}{e} \geq \frac{\gamma}{1+\gamma} \\ -\gamma \cdot \log \frac{\gamma}{1+\gamma} & \text{else} \end{cases}.$$

Note that b is a continuous function. Further, for $\gamma = 1$ the bound is $b(1) = \ln(2)$ and that for large enough γ , namely $\gamma \gtrsim 1.32$, this bound is dominated by the upper bound of 0.745 [Correa et al., 2017].

5 Superlinear Number of Samples

Our final result is that it is in fact possible to get arbitrarily close to the optimal approximation guarantee of a stopping algorithm that knows the distribution [Correa et al., 2017], if we have access to $O(n^2)$ samples from the distribution. We provide details and the proof in the appendix.

Theorem 5.1. *Let X_1, \dots, X_n be i.i.d. random variables drawn from an unknown distribution F . Then for every $\varepsilon > 0$ and all $n \geq n_\varepsilon$ there exists an algorithm for choosing a stopping time τ that uses $O(n^2)$ samples from the same distribution with*

$$\mathbb{E}[X_\tau] \geq (0.745 - \varepsilon) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

In the appendix we also provide evidence that any algorithm that achieves this bound with $o(n^2)$ samples would have to use very different techniques.

Acknowledgements

We are grateful to Fábio Botler for suggesting the use of Ramsey's theorem in the proof of Theorem 3.1, and to the anonymous referees for many valuable comments on an earlier version of the paper.

APPENDIX

A Proofs from Subsection 3.1

Proof of Theorem 3.1. Let τ be the stopping time corresponding to the optimal stopping rule for the secretary problem, which rejects a certain fraction of the random variables and uses their maximum as a threshold for the remaining ones. Since X_1, X_2, \dots, X_n are drawn independently from the same distribution, we can assume that their realizations are obtained by independently drawing n values from the distribution and then ordering them according to a random permutation π . Denoting the density of the distribution from which X_1, \dots, X_n are drawn by f ,

$$\begin{aligned} \mathbb{E}[X_\tau] &= \int_0^\infty \cdots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \mathbb{E}_\pi[v_{\pi(\tau)}] \, dv_1 \cdots dv_n \\ &\geq \int_0^\infty \cdots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \Pr_\pi[v_{\pi(\tau)} = \max\{v_1, \dots, v_n\}] \cdot \max\{v_1, \dots, v_n\} \, dv_1 \cdots dv_n \\ &\geq \frac{1}{e} \cdot \int_0^\infty \cdots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \max\{v_1, \dots, v_n\} \, dv_1 \cdots dv_n \\ &= \frac{1}{e} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}], \end{aligned}$$

where the second inequality holds because the values v_1, \dots, v_n have been randomly ordered and τ is thus guaranteed to select $\max\{v_1, \dots, v_n\}$ with probability at least $1/e$ for any realization [Ferguson, 1989]. This proves the claim. \square

B Proofs from Subsection 3.2

Proof of Lemma 3.4. For $i \in [n]$, let \sim_i be the equivalence relation on \mathbb{R}_+^i such that $(v_1, \dots, v_i) \sim_i (w_1, \dots, w_i)$ if v_1, \dots, v_{i-1} is a permutation of w_1, \dots, w_{i-1} and $v_i = w_i$. Note that a stopping time τ is order-oblivious if and only if for all $i \in [n]$ and $v_1, \dots, v_i, w_1, \dots, w_i \in \mathbb{R}_+$ it holds that $p_i^\tau(v_1, \dots, v_i) = p_i^\tau(w_1, \dots, w_i)$ whenever $(v_1, \dots, v_i) \sim_i (w_1, \dots, w_i)$. We will refer to the equivalence classes of \sim_i as *states*, and will say that a stopping time τ *arrives at* $s \in \mathbb{R}_+^i / \sim_i$ in the event that $\tau \geq i$ and $X_1 = v_1, \dots, X_{i-1} = v_{i-1}$ where $[v_1, \dots, v_i]_{\sim_i} = s$.

Let τ be an arbitrary stopping time, and define a stopping time σ such that $p_1^\sigma(v_1) = p_1^\tau(v_1)$ and for all $i \in \{2, \dots, n\}$ and $v_1, \dots, v_i \in \mathbb{R}_+$ with $\Pr[\tau \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] > 0$,

$$p_i^\sigma(v_1, \dots, v_i) = \Pr[\tau = i \mid \tau \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}].$$

Since $[v_1, \dots, v_i]_{\sim_i}$ is invariant under permutations of the sequence v_1, \dots, v_{i-1} , σ is indeed order-oblivious. It remains to be shown that σ provides guarantee α .

As an intermediate step we show by induction that for all $i \in [n]$ and $s \in \mathbb{R}_+^i / \sim_i$,

$$\Pr[\tau \text{ arrives at } s] = \Pr[\sigma \text{ arrives at } s]. \quad (2)$$

This holds trivially for $i = 1$, so we assume that it holds for $i = k - 1 \geq 1$ and show then that it holds for $i = k$. Indeed, for any $v_1, \dots, v_k \in \mathbb{R}_+$ and $s = [v_1, \dots, v_k]_{\sim_k}$, writing $v_{-j} = (v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_{k-1})$ for the $k - 2$ dimensional vector in which we leave out v_j ,

$$\Pr[\tau \text{ arrives at } s] = \sum_{j=1}^{k-1} \Pr[\tau \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[\tau \neq i \mid \tau \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[X_k = v_k]$$

$$\begin{aligned}
&= \sum_{j=1}^{k-1} \Pr[\sigma \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[\sigma \neq i \mid \sigma \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[X_k = v_k] \\
&= \Pr[\sigma \text{ arrives at } s],
\end{aligned}$$

where the first and last equalities hold by definition of \sim_{k-1} and the second equality by the induction hypothesis and by definition of σ .

We now claim that

$$\begin{aligned}
\mathbb{E}[X_\tau] &= \sum_{i=1}^n \mathbb{E}[X_i \mid \tau = i] \cdot \Pr[\tau = i] \\
&= \sum_{i=1}^n \int_0^\infty \cdots \int_0^\infty \prod_{j=1}^i f(v_j) \cdot v_i \\
&\quad \cdot \frac{1}{(i-1)!} \cdot \sum_{\pi \in \mathcal{S}_{i-1}} \Pr[\tau = i \mid X_1 = v_{\pi(1)}, \dots, X_{i-1} = v_{\pi(i-1)}, X_i = v_i] \, dv_1 \dots dv_i \\
&= \sum_{i=1}^n \int_0^\infty \cdots \int_0^\infty \prod_{j=1}^i f(v_j) \cdot v_i \cdot \Pr[\tau = i \mid \tau \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] \\
&\quad \cdot \Pr[\tau \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] \, dv_1 \dots dv_i \\
&= \sum_{i=1}^n \int_0^\infty \cdots \int_0^\infty \prod_{j=1}^i f(v_j) \cdot v_i \cdot p_i^\sigma(v_1, \dots, v_i) \cdot \Pr[\sigma \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] \, dv_1 \dots dv_i \\
&= \sum_{i=1}^n \mathbb{E}[X_i \mid \sigma = i] \cdot \Pr[\sigma = i] = \mathbb{E}[X_\sigma].
\end{aligned}$$

Indeed, the second equality can be seen to hold by imagining that X_1, \dots, X_i are drawn by first drawing i values independently and then permuting the first $i-1$ of these values uniformly at random. The fourth equality holds by definition of σ and by (2). This completes the proof. \square

C Proofs from Subsection 3.4

Proof of Corollary 3.6. For $\delta > 0$, choose $\gamma > 0$ such that $(1 + \gamma)/e \leq 1/e + \delta/2$ and $\gamma/(1 + \gamma) \leq 1/e$. By Theorem 4.5, there exists an n_1 such that for all $n \geq n_1$ and all $(\gamma n, n)$ -stopping rules with associated stopping time τ there exists a distribution F , not known to the stopping rule, with the following property. When X_1, \dots, X_n are i.i.d. random variables drawn from F , we have

$$\mathbb{E}[X_\tau] \leq \left(\frac{1 + \gamma}{e} + \frac{\delta}{2} \right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}] \leq \left(\frac{1}{e} + \delta \right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where the second inequality follows by our choice of γ .

Now let n_0 be such that $f(n) \leq \gamma n$ for all $n \geq n_0$. As every $(f(n), n)$ -stopping rule can be interpreted as a $(\gamma n, n)$ -stopping rule when $n \geq n_0$, the above bound for $(\gamma n, n)$ -stopping rules applies to $(f(n), n)$ -stopping rules as well when $n \geq \max\{n_0, n_1\}$. This proves the claim. \square

D A 1/2-approximation with $n - 1$ samples

In this appendix we formalize the discussion in Section 4.1. We show that if the stopping rule has access to $n - 1$ samples, then we can simply take the maximum of these samples as a single, non-adaptive threshold for all random variables to obtain a factor 1/2-approximation.

Theorem D.1. Let X_1, X_2, \dots, X_n be i.i.d. draws from an unknown distribution F . Then there exists a $(n-1, n)$ -stopping-rule with stopping time τ such that

$$\mathbb{E}[X_\tau] \geq \frac{1}{2} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

To prove Theorem D.1 we will analyze a slight variation of the algorithm described above, Algorithm 2, which only uses the maximum of $n-1$ samples as a threshold for the first $n-1$ random variables and stops on the n th random variable with certainty. The advantage of this is that it becomes even clearer when and why our analysis is tight.

Algorithm 2: Single threshold algorithm

Data: Sequence of i.i.d. random variables X_1, \dots, X_n sampled from an unknown distribution F ,
sample access to F

Result: Stopping time τ

$\tau \leftarrow n$

$S_1, \dots, S_{n-1} \leftarrow n-1$ samples from F

for $t = 1, \dots, n-1$ **do**

if $X_t \geq \max\{S_1, \dots, S_{n-1}\}$ **then**
 $\tau \leftarrow t$
 Break

return τ

Proof of Theorem D.1. The expected value achieved by Algorithm 2 is the sum over all time steps $i = 1, \dots, n$ of the product of the probability of stopping at this time step and the expected value of the random variable conditioned on being above the threshold

$$\begin{aligned} \mathbb{E}[X_\tau] &= \sum_{i=1}^{n-1} \left(\mathbb{E}[X_i \mid \tau = i] \cdot \Pr[\tau = i] \right) + \mathbb{E}[X_n] \cdot \Pr[\tau = n] \\ &\geq \sum_{i=1}^{n-1} \left(\mathbb{E}[X_i \mid \tau = i] \cdot \Pr[\tau = i] \right). \end{aligned} \quad (3)$$

We stop at time step i if the maximum among the $n-1$ samples and the first i random variables happens to be the i th random variable, and if, conditioned on this, the second maximum is among the $n-1$ samples and not the other $i-1$ random variables. Hence,

$$\Pr[\tau = i] = \frac{1}{n-1+i} \cdot \frac{n-1}{n-2+i}$$

Summing this over all i from 1 to $n-1$ shows that the probability of stopping at one of the first $n-1$ random variables is precisely

$$\sum_{i=1}^{n-1} \Pr[\tau = i] = \sum_{i=1}^{n-1} \frac{1}{n-1+i} \cdot \frac{n-1}{n-2+i} = \frac{1}{2}. \quad (4)$$

We conclude the proof by showing that for all $i = 1, \dots, n-1$ the conditional expectation $\mathbb{E}[X_i \mid \tau = i]$ is at least $\mathbb{E}[\max\{X_1, \dots, X_n\}]$. Let $T = \max\{S_1, \dots, S_n\}$. The algorithm stops at time step i if $X_i \geq T > \max\{X_1, \dots, X_{i-1}\}$. So under this event X_i is the maximum of $n-1+i$ random variables. And so

$$\mathbb{E}[X_i \mid \tau = i] = \mathbb{E}[\max \text{ of } n-1+i \text{ i.i.d. RVs}] \geq \mathbb{E}[\max\{X_1, \dots, X_n\}]. \quad (5)$$

Substituting (4) and (5) into (3) completes the proof. \square

As argued in the proof of Theorem D.1 the probability that Algorithm 2 stops on one of the first $n - 1$ variables is precisely $1/2$. The two potentially lossy steps are that we dropped the contribution from the final random variable, and that we lower bounded the contribution from each of the first $n - 1$ random variables by $\mathbb{E}[\max\{X_1, \dots, X_n\}]$.

It turns out that both of the potentially lossy steps are in fact lossless in the limit as $n \rightarrow \infty$ if F is the exponential distribution.

Proposition D.2. *Let X_1, \dots, X_n be drawn independently from the exponential distribution $F = 1 - e^{-x}$. Then for the stopping time τ determined by Algorithm 2,*

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[X_\tau]}{\mathbb{E}[\max\{X_1, \dots, X_n\}]} = \frac{1}{2}.$$

Proof. It is not super straightforward to compute, but a well-known fact that the maximum of n independent, exponentially distributed random variables X_1, \dots, X_n is equal to the n th Harmonic number. That is,

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] = H_n$$

As argued in the proof of Theorem D.1 we can express the expected value obtained by Algorithm 2 as follows

$$\mathbb{E}[X_\tau] = \sum_{i=1}^{n-1} \left(H_{n-1+i} \cdot \frac{1}{n-1+i} \cdot \frac{n-1}{n-2+i} \right) + \frac{1}{2}.$$

Tedious calculations allow to express the expected value via the digamma function $\psi^{(0)}$ and the Euler-Mascheroni constant γ as follows

$$\mathbb{E}[X_\tau] = \psi^{(0)}(n) - \frac{1}{2}H_{2n-2} + \gamma + 1.$$

This can then be used to show that

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[X_\tau]}{\mathbb{E}[\max\{X_1, \dots, X_n\}]} = \lim_{n \rightarrow \infty} \frac{\psi^{(0)}(n) - \frac{1}{2}H_{2n-2} + \gamma + 1}{H_n} = \frac{1}{2},$$

which proves the claim. □

We conclude by showing that our analysis of Algorithm 2 is tight, and indeed that any stopping rule that beats the bound of $1/2$ has to use a different approach.

Proposition D.3. *Let $\varepsilon > 0$. Then there exists a distribution F and $n \in \mathbb{N}$ such that for random variables X_1, \dots, X_n drawn independently from this distribution and any stopping time τ induced by setting the same threshold to each of the random variables*

$$\mathbb{E}[X_\tau] \leq \left(\frac{1}{2} + \varepsilon \right) \cdot \mathbb{E} \left[\max_t X_t \right].$$

Proof. Take n copies of the random variable X that is

$$X_n = \begin{cases} n & \text{with probability } 1/n^2, \text{ and} \\ 1 & \text{with probability } 1 - 1/n^2. \end{cases}$$

Then $\mathbb{E}[\max\{X_1, \dots, X_n\}] = (1 - (1 - 1/n^2)^n) \cdot n + (1 - 1/n^2)^n \cdot 1$ which tends to 2 as $n \rightarrow \infty$. The only two sensible thresholds are $T = n$ or $T = 1$. For $T = n$ we obtain $\mathbb{E}[X_\tau] = (1 - (1 - 1/n^2)^n) \cdot n$ which tends to 1 as $n \rightarrow \infty$. For $T = 1$ we obtain $\mathbb{E}[X_\tau] = 1$. □

E Proofs from Subsection 4.3

Proof of Proposition 4.3. For each $i \in [n]$, let

$$X_i = \begin{cases} \frac{\sqrt{n}}{e-2} & \text{w.p. } \frac{1}{n^{3/2}}, \\ 1 & \text{w.p. } \frac{1}{\sqrt{n}}, \\ 0 & \text{otherwise.} \end{cases}$$

It is not difficult to see that, when n is large enough,

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] \geq \frac{1}{e-2} + 1 - \varepsilon.$$

On the other hand, we can restrict our attention to stopping rules that always accept a value of $\frac{\sqrt{n}}{e-2}$ and never accept a value of 0. A stopping rule with the property that the probability of stopping conditioned on reaching the i th random variable is independent of i , is thus characterized by a probability α with which it accepts a value of 1. Denote by τ_α the stopping rule that accepts a value of 1 with probability α . Then, when n is large enough,

$$\mathbb{E}[X_{\tau_\alpha}] = \mathbb{E}[X_{\tau_\alpha} \mid X_{\tau_\alpha} > 0] \cdot \Pr[X_{\tau_\alpha} > 0],$$

where

$$\begin{aligned} \mathbb{E}[X_{\tau_\alpha} \mid X_{\tau_\alpha} > 0] &= \frac{1/n^{3/2}}{1/n^{3/2} + \alpha/\sqrt{n}} \cdot \frac{\sqrt{n}}{e-2} + \frac{\alpha/\sqrt{n}}{1/n^{3/2} + \alpha/\sqrt{n}} \cdot 1 \\ &= \frac{1/(n(e-2)) + \alpha/\sqrt{n}}{1/n^{3/2} + \alpha/\sqrt{n}} \end{aligned}$$

and

$$\begin{aligned} \Pr[X_{\tau_\alpha} > 0] &= 1 - \left(1 - \frac{1}{n^{3/2}} - \frac{\alpha}{\sqrt{n}}\right)^n \\ &\leq 1 - e^{-(\alpha\sqrt{n} + 1/\sqrt{n})} + \varepsilon. \end{aligned}$$

Thus

$$\mathbb{E}[X_{\tau_\alpha}] \leq \frac{1/(e-2) + \alpha\sqrt{n}}{1/\sqrt{n} + \alpha\sqrt{n}} \cdot (1 - e^{-(\alpha\sqrt{n} + 1/\sqrt{n})} + \varepsilon).$$

Behavior of the expression on the right-hand side is determined by the value of $\alpha\sqrt{n}$. If $\alpha\sqrt{n}$ increases in n then $\mathbb{E}[X_{\tau_\alpha}]$ tends to 1. If $\alpha\sqrt{n}$ decreases in n then $e^{-(\alpha\sqrt{n} + 1/\sqrt{n})} \geq 1 - (\alpha\sqrt{n} + 1/\sqrt{n}) + \varepsilon$, so that $\mathbb{E}[X_{\tau_\alpha}]$ tends to $1/(e-2)$. The case where $\alpha\sqrt{n}$ is constant can finally be solved by considering the maximum of $((1/(e-2) + x)/x) \cdot (1 - e^{-x})$, which occurs at $x = 1$.

In summary

$$\frac{\mathbb{E}[X_{\tau_\alpha}]}{\mathbb{E}[\max\{X_1, \dots, X_n\}]} \leq \left(1 - \frac{1}{e} + \varepsilon\right),$$

which shows the claim. □

Proof of Proposition 4.4. Take an instance with $X_i = 1$ with probability ε/n and 0 otherwise (actually we take values that are small perturbations of 1 and 0 to avoid ties and make the distribution continuous). For this distribution, all we care about are situations in which among all the samples and values there is exactly one random variable with value 1: indeed if there are only 0's the instance is irrelevant while having more than one value 1 is extremely unlikely. Furthermore, relevant instances are those in which this only 1 value appears in X_1, \dots, X_n and not in the samples, say $X_k = 1$.

Now it is easy to bound the probability that an algorithm satisfying (i) and (ii) stops before time k . Say α_i is the probability of stopping at X_i conditional on reaching it. It follows that $\Pr[\tau < k] = \alpha_1 + (1 - \alpha_1)(\alpha_2 + (1 - \alpha_2)(\dots))$. This is linear in α_1 and the coefficient is non-negative (it is 1 minus some probability). So this is minimized at $\alpha_1 = 1/n$. Of course this holds for all i . So,

$$\Pr[\tau < k] \geq \sum_{j=0}^{k-1} \left(1 - \frac{1}{n}\right)^j \cdot \frac{1}{n} = 1 - \left(1 - \frac{1}{n}\right)^k.$$

Now notice that as the instance is i.i.d. the value of k is uniform in $\{1, \dots, n\}$, so that we can lower bound the probability that an algorithm satisfying (i) and (ii) stops before seeing the value 1 as follows:

$$\begin{aligned} \Pr[\text{stop before the value 1}] &= \sum_{k=1}^n \Pr[\text{stop before the value 1} \mid \text{the value 1 appears at position } k] \cdot \frac{1}{n} \\ &= \frac{1}{n} \cdot \sum_{k=1}^n \Pr[\tau < k] \\ &\geq \frac{1}{n} \cdot \sum_{k=1}^n \left(1 - \left(1 - \frac{1}{n}\right)^k\right) = \left(1 - \frac{1}{n}\right)^{n+1} + \frac{1}{n} \approx 1/e. \end{aligned}$$

Therefore, the algorithm misses the value 1 with probability $1/e$, so at best it can get it with probability $1 - 1/e$, and thus this is the best hope for an approximation ratio. \square

F Proofs from Subsection 4.4

Proof of Theorem 4.5. Consider $n \in \mathbb{N}$ such that $\gamma n \in \mathbb{N}$ and a $(\gamma n, n)$ -stopping rule \mathbf{r} . Let τ be its stopping time. Set $\varepsilon = 1/n^2$. By Lemma 3.3 there exists an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{r} is ε -value-oblivious. Let $v_1, \dots, v_{n^3}, u \in V$ be pairwise distinct such that $u \geq n^3 \max\{v_1, \dots, v_{n^3}\}$. For each $i \in [n]$, let

$$X_i = \begin{cases} v_1 & \text{w.p. } \frac{1}{n^3} \cdot \left(1 - \frac{1}{n^2}\right) \\ \vdots & \\ v_{n^3} & \text{w.p. } \frac{1}{n^3} \cdot \left(1 - \frac{1}{n^2}\right) \\ u & \text{w.p. } \frac{1}{n^2} \end{cases}$$

We proceed to bound $\mathbb{E}[\max\{X_1, \dots, X_n\}]$ from below and $\mathbb{E}[X_\tau]$ from above. For $i \in [(1 + \gamma) \cdot n]$, let

$$R_i = \begin{cases} S_i & \text{if } i \leq \gamma n \\ X_{i-n+1} & \text{otherwise.} \end{cases}$$

Let $\tau' = \tau + \gamma n$ be the stopping time of \mathbf{r} on $R_1, \dots, R_{(1+\gamma)n}$. Let $X_{(i)}$ denote i th order statistic of X_1, \dots, X_n , such that $X_{(n)} = \max\{X_1, \dots, X_n\}$. Then

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] \geq \Pr[X_{(n)} = u] \cdot u = \frac{1 - o(1)}{n} \cdot u.$$

On the other hand,

$$\begin{aligned}
\mathbb{E}[X_\tau] &= \Pr[R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} \neq u] \cdot \mathbb{E}[R_\tau \mid R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} \neq u] \\
&\quad + \Pr[R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} = u] \cdot \mathbb{E}[R_{\tau'} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} = u] \\
&\quad + \Pr[R_{((1+\gamma)\cdot n)} \neq u] \cdot \mathbb{E}[R_{\tau'} \mid R_{((1+\gamma)\cdot n)} \neq u] \\
&\leq \frac{1+\gamma}{n} \cdot \left(\Pr[R_{\tau'} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} \neq u] \cdot u \right. \\
&\quad \left. + \Pr[R_{\tau'} \neq R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} \neq u] \cdot O(n^{-3}) \cdot u \right) \\
&\quad + O(n^{-2}) \cdot u + 1 \cdot O(n^{-3}) \cdot u \\
&\leq \frac{1+\gamma+o(1)}{n} \cdot \Pr[R_{\tau'} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_{((1+\gamma)\cdot n-1)} \neq u] \cdot u \\
&\leq \frac{1+\gamma+o(1)}{n} \cdot \Pr[R_{\tau'} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_1, \dots, R_{(1+\gamma)\cdot n} \text{ are distinct}] \cdot u.
\end{aligned}$$

Let

$$p = \Pr[R_{\tau'} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_1, \dots, R_{(1+\gamma)\cdot n} \text{ are distinct}].$$

To complete the proof, we bound p . To this end, let $w_1, \dots, w_{(1+\gamma)\cdot n} \in \{v_1, \dots, v_{n^3}, u\}$ such that $w_1 < \dots < w_{(1+\gamma)\cdot n} = u$. Under the condition that $\{R_1, \dots, R_{(1+\gamma)\cdot n}\} = \{w_1, \dots, w_{(1+\gamma)\cdot n}\}$, the values appear in a random order. Moreover, r is ε -value-oblivious on V . For each $i \in [(1+\gamma)\cdot n]$ and up to an error probability of $((1+\gamma)\cdot n) \cdot 2 \cdot \varepsilon = O(1/n)$, it must therefore select R_i with the same probability q_i if $R_i > \max\{R_1, \dots, R_{i-1}\}$. But since $R_1, \dots, R_{\gamma n}$ are in fact samples, $q_i = 0$ for all $i = n, \dots, (1+\gamma)\cdot n$. We are thus faced with an instance of the secretary problem with $(1+\gamma)\cdot n$ values under the additional constraint that the first γn values have to be rejected.

The optimal stopping rule for this problem is known to set, for some $x \in [0, 1]$, $q_i = 0$ for all $i < x \cdot (1+\gamma)\cdot n$ and $q_i = 1$ for all $i \geq x \cdot (1+\gamma)\cdot n$ [Gilbert and Mosteller, 1966]. Then $p = -x \cdot \log x + o(1)$, which subject to $x \geq \gamma/(1+\gamma)$ is maximized for

$$x = \max \left\{ \frac{1}{e}, \frac{\gamma}{1+\gamma} \right\}$$

and thus $p \leq b(\gamma)/(1+\gamma) + o(1)$ and $\mathbb{E}[X_\tau] < (b(\gamma) + o(1)) \cdot u$ and where

$$b(\gamma) = \begin{cases} \frac{1+\gamma}{e} & \text{if } \frac{1}{e} \geq \frac{\gamma}{1+\gamma} \\ -\gamma \cdot \log \frac{\gamma}{1+\gamma} & \text{else} \end{cases}.$$

Since this is true for every sequence of values $w_1, \dots, w_{(1+\gamma)\cdot n}$ as above, we are done. \square

G Details for Section 5 and Proof of Theorem 5.1

We first recall that the optimal algorithm for known distributions computes a decreasing sequence $x_i = y(i/n)^{1/(n-1)}$ for $i \in [n]$, where y is the unique solution (which turns out to be decreasing and convex) to the following ordinary differential equation

$$y' = y(\ln(y) - 1) - (\beta - 1) \quad \text{and} \quad y(0) = 1,$$

where $\beta \approx 1.3414 \approx 1/0.745$. Then conditional on reaching random variable X_i , the algorithm accepts it with probability essentially equal to $\varepsilon_i = 1 - x_i$; where for consistency we have that $0 = \varepsilon_0 < \varepsilon_1 <$

$\dots < \varepsilon_{n-1} < \varepsilon_n = 1$. For $q \in [0, 1]$ let $R(q) = \int_0^q F^{-1}(1 - \theta) d\theta$. Let τ be the stopping time implied by the algorithm. It can then be shown that

$$\mathbb{E}[X_\tau] = \sum_{i=1}^n 0.745 \cdot n \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2} R(q) dq = 0.745 \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where the i th term in the sum is the contribution of X_i to the expectation.

Proof sketch for Theorem 5.1. The algorithm that achieves the claimed bound starts by skipping some random variables until the acceptance probability of the optimal algorithm ε_i becomes sufficiently large, say δ/n , say this happens at step k . Afterwards, it uses the empirical distribution function of the samples to estimate the quantiles $\varepsilon_{k+1}, \dots, \varepsilon_n$ used by the optimal algorithm that knows the distribution on the remaining random variables. The algorithm then accepts random variable X_i conditional on reaching it with probability $\tilde{\varepsilon}_i$, where $\tilde{\varepsilon}_i$ is its estimate of ε_i . The reason why we skip the first few random variables is because the initial acceptance probability of the optimal algorithm is of the order of $1/n^2$, therefore with n^2 sample we cannot get a reliable estimate of the corresponding threshold.

To see that the algorithm satisfies the claimed property we first observe that by skipping the first few random variables, until the acceptance probability is δ/n we only lose a small revenue. Indeed, the revenue of the algorithm until the acceptance probability becomes δ/n is given by:

$$n \cdot 0.745 \int_0^{\delta/n} (n-1)(1-q)^{n-2} R(q) dq.$$

Since $R(q)$ is monotone, this revenue is at most

$$n \cdot 0.745 \int_0^{\delta/n} (n-1)(1-q)^{n-2} dq \cdot R(\delta/n) \leq n \cdot 0.745(1 - e^{-\delta}) \cdot R(\delta/n) \leq n \cdot 0.745\delta \cdot R(\delta/n).$$

On the other hand, between this time and the end the algorithm's expected revenue is

$$n \cdot 0.745 \int_{\delta/n}^1 (n-1)(1-q)^{n-2} R(q) dq \geq n \cdot 0.745 e^{-\delta} \cdot R(\delta/n) \geq n \cdot 0.745(1 - \delta) \cdot R(\delta/n).$$

Therefore the ratio between the two is $\delta/(1 - \delta)$, so that by picking a small value δ , the loss can be made arbitrarily small.

The next observation is that because we have $O(n^2)$ samples we can use the Dvoretzky-Kiefer-Wolfowitz inequality [1956] to argue that with probability at least $1 - \alpha$ all quantiles of the empirical distribution function will be within a β/n band of the actual quantiles, with β an arbitrarily small constant.

Now conditioned on the fact that all our estimated quantiles lie within the respective error band the probabilities with which our algorithm stops on each of the random variables X_{k+1}, \dots, X_n is arbitrarily close to the corresponding acceptance probabilities in the optimal algorithm that knows the distribution. Indeed as the error within each step is at most δ/n for the first k random variables and β/n for the random variables starting from k , the cumulative error until step i is at most $1 - (1 - \max\{\delta, \beta\}/n)^i \approx \max\{\delta, \beta\}$, which by the choice of δ can be made arbitrarily small. The latter shows that the distribution of the stopping time of our algorithm and that of the optimal algorithm are essentially the same.

The last ingredient in the proof is to notice that conditional on stopping at a given time both algorithms get roughly the same amount. This is quite clear since in general for a random variable X and two thresholds τ_1 and τ_2 such that $F(\tau_1)$ and $F(\tau_2)$ are close then $\mathbb{E}[X | X > \tau_1]$ is close to $\mathbb{E}[X | X > \tau_2]$.

In summary our algorithm skips the first few random variables by losing only a small fraction of the revenue, then stops essentially at the same (random) time than the optimal algorithm, and finally conditional on stopping at time i it obtains essentially the same reward as the optimal algorithm. \square

H Achieving 0.745 with $o(n^2)$ Samples

We conclude with an argument that suggests that any approach that achieves the optimal 0.745 approximation with $o(n^2)$ samples would have to go through very different techniques. To this end we will show that even if the random variables X_1, \dots, X_n were uniform on $[0, 1]$ in order to get concentration of the median around its expectation within a $O(1/n)$ band at least $\Omega(n^2)$ samples are needed.

Proposition H.1. *Let X denote the median of $f(n)$ samples $X_1, \dots, X_{f(n)}$ drawn independently from the uniform distribution on $[0, 1]$. Then, for any constant $\varepsilon > 0$, $f(n)$ must be in the order of $\Omega(n^2)$ to have*

$$\Pr \left[|X - \mathbb{E}[X]| \leq \frac{1}{n} \right] \leq \varepsilon.$$

Proof of Proposition H.1. The k th smallest of n samples from a uniform distribution follows a $\text{Beta}(n - k + 1, k)$ distribution. So if we are interested in the median of n^2 samples, then for n large, the distribution of the median is well approximated by $X \sim \text{Beta}(n^2/2, n^2/2)$.

The expectation of a random variable drawn from $\text{Beta}(\alpha, \beta)$ is $\alpha/(\alpha + \beta)$. So when $\alpha = \beta$ as in our case this is simply $1/2$. The variance of such a random variable is $\alpha \cdot \beta / [(\alpha + \beta)^2 \cdot (\alpha + \beta + 1)]$. So for $\alpha = \beta = n^2/2$ it is $1/[4(n^2 + 1)]$. Now we want to show concentration within a $1/n$ band around the expectation, which is $1/2$. For simplicity, we will look at the one sided error only. So we seek to bound

$$\Pr \left[X \geq \frac{1}{2} + \frac{1}{n} \right].$$

To compute this probability, we will use that for $\alpha = \beta$ large, we can approximate the Beta distribution with a normal distribution. More formally, for $\alpha = \beta = n^2/2$ large the random variable $Y = 2 \cdot \sqrt{n^2 + 1} \cdot (X - 1/2)$ has probability density function

$$f_Y(y) = \frac{1}{\sqrt{\pi}} \cdot e^{-\frac{y^2}{2}}.$$

With this we obtain

$$\begin{aligned} \Pr \left[X \geq \frac{1}{2} + \frac{1}{n} \right] &\approx \Pr \left[Y \geq 2 \cdot \sqrt{n^2 - 1} \cdot \frac{1}{n} \right] \\ &= \frac{1}{2} - \frac{1}{2} \cdot \text{erf} \left(2 \cdot \sqrt{n^2 - 1} \cdot \frac{1}{n} \cdot \frac{1}{\sqrt{2}} \right). \end{aligned}$$

Now the argument of the Erlang function tends to $\sqrt{2}$ as n tends to ∞ , and $\text{erf}(\sqrt{2}) \approx 0.954$. So the probability that the one sided error is at most $1/n$ is close to 1, and we can make it arbitrarily close to 1 by using $c \cdot n^2$ samples instead.

If on the other hand we used $o(n^2)$ samples, the argument of the Erlang function would tend to zero and hence the probability would not vanish as desired. \square

References

- M. Abolhassani, S. Ehsani, H. Esfandiari, M. Hajiaghayi, R. D. Kleinberg, and B. Lucier. Beating $1 - 1/e$ for ordered prophets. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, pages 61–71, 2017.
- S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

- P. D. Azar, R. Kleinberg, and S. M. Weinberg. Prophet inequalities with limited information. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1358–1377, 2014.
- M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer. Posting prices with unknown distributions. In *Proceedings of the 1st Symposium on Innovations in Computer Science*, pages 166–178, 2011.
- B. A. Berezovskiy and A. V. Gnedin. *The Best Choice Problem (In Russian)*. Nauka, Moscow, 1984.
- S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 311–320, 2010.
- S. Chawla, J. Miller, and Y. Teng. Posted pricing for online resource allocation: Beyond subadditive values. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2019. Forthcoming.
- J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 169–186, 2017.
- A. V. den Boer. Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1):1 – 18, 2015.
- P. Dütting and R. Kleinberg. Polymatroid prophet inequalities. In *Proceedings of the 23rd European Symposium on Algorithms*, pages 437–449, 2015.
- P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science*, pages 540–551, 2017.
- A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Annals of Mathematical Statistics*, 27:642–669, 1956.
- S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–714, 2018.
- M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, pages 123–135, 2015.
- M. Feldman, O. Svensson, and R. Zenklusen. Online contention resolution schemes. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms*, pages 1014–1033, 2016.
- T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
- J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.
- A. Goldenshluger and A. Zeevi. Optimal stopping of a random sequence with unknown distribution. Working paper, 2017.
- M. Hajiaghayi, R. Kleinberg, and T. W. Sandholm. Automated mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 58–65, 2007.

- T. Hill and R. P. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics*, 125:191–207, 1992.
- T. P. Hill and R. P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability*, 10:336–345, 1982.
- R. P. Kertz. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis*, 19:88–112, 1986.
- R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th ACM Symposium on Theory of Computing Conference*, pages 123–136, 2012.
- U. Krengel and L. Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83:745–747, 1977.
- U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Probability Related Topics*, 4:197–266, 1978.
- F. P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 30:264–286, 1930.
- A. Rubinstein. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the 48th ACM Symposium on Theory of Computing*, 2016. 324–332.
- A. Rubinstein and S. Singla. Combinatorial prophet inequalities. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687, 2017.
- E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12:1213–1216, 1984.