

Prophet Inequalities for I.I.D. Random Variables from an Unknown Distribution

JOSÉ CORREA, Universidad de Chile, Chile

PAUL DÜTTING, London School of Economics, UK

FELIX FISCHER, Queen Mary University of London, UK

KEVIN SCHEWIOR, École Normale Supérieure, France and Technische Universität München, Germany

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 a long time the best known bound for the problem had been $\alpha \geq 1 - 1/e \approx 0.632$, but quite recently a tight bound of $\alpha \approx 0.745$ was obtained. The case where F is unknown, such that the decision whether $\tau = t$ may depend only on the values of the random variables X_1, \dots, X_t , is equally well motivated 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, where an arbitrary set of values arrive in random order and the goal is to maximize the probability of selecting the largest value. We show that this bound is in fact tight. We then investigate the case where the stopping time may additionally depend on a limited number of samples from F , and show that even with $o(n)$ samples $\alpha \leq 1/e$. On the other hand, n samples allow for a significant improvement, 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 $\alpha \geq 0.745 - \epsilon$ for any $\epsilon > 0$.

CCS Concepts: • **Theory of computation** → **Online algorithms**; **Computational pricing and auctions**; • **Applied computing** → **Economics**.

Additional Key Words and Phrases: posted pricing, prophet inequalities

ACM Reference Format:

José Correa, Paul Dütting, Felix Fischer, and Kevin Schewior. 2019. Prophet Inequalities for I.I.D. Random Variables from an Unknown Distribution. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19), June 24–28, 2019, Phoenix, AZ, USA*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3328526.3329627>

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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329627>

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., 17]. 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 [5, 16]. When numbers are drawn independently from a single known distribution a better (and tight) guarantee of around 0.58 can be obtained [17].

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 [24, 25] and Samuel-Cahn [29]. For the latter a lower bound of $1 - 1/e \approx 0.632$ due to Hill and Kertz [21] was improved only very recently, first to 0.738 by Abolhassani et al. [1] and then to 0.745 by Correa et al. [8]. The bound of Correa et al. is in fact known to be tight due to a matching upper bound of Hill and Kertz [21] and Kertz [22].

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., [3]), 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 (Proposition 3.1 in Section 3). 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.

We show that the straightforward guarantee of $1/e$ is in fact best possible in the prophet setting (Theorem 3.2 in Section 3). 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 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




Fig. 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 $\Theta(n^2)$ samples are tight. With the exception of the upper bound of approximately 0.745, all results are new to this paper.

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 $\epsilon > 0$, a subsequence of values contained in the interval $[q - \epsilon, q + \epsilon]$. 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 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 [26] 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 upper bound construction 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. We show that a simple but subtle stopping rule achieves an approximation factor of $1 - 1/e \approx 0.632$ with $n - 1$ samples (Theorem 4.1 in Section 4). 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 $1 - 1/e$. By a straightforward extension (Corollary 4.5), we obtain a lower bound of $\frac{1+\alpha}{2} \cdot (1 - 1/e)$ for $\alpha < 1$.

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 $\gamma \geq 0$ (Theorem 4.6 in Section 4). 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 (Theorem 5.1 in Section 5). 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 [12] 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 [20]. Starting from work of Hajiaghayi et al. [19] prophet inequalities, and in particular extensions to richer feasibility domains, have seen a surge of interest in theoretical computer science (e.g., [2, 6, 7, 10, 11, 13–15, 23, 27, 28]).

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. [3], which focuses on richer feasibility structures such as matching constraints and matroids, and work by Babai et al. [4], 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. In concurrent work Wang [30] considers the case of unknown, non-identical distributions and shows how to obtain a factor 2 approximation with n samples.

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 [18] and the recent survey of den Boer [9]).

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 in this section, 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 [16] 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 full version of this paper.

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 3.3. 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 3.4 (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 3.4 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 3.4 there exists a stopping rule \mathbf{r} with performance guarantee $1/e + \delta$ and

an infinite set $V \subseteq \mathbb{N}$ on which r is ε -value-oblivious. Denote by τ the stopping time of 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

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

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 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{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{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 r and \hat{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 $s_1, \dots, s_n \in V$,

$$\begin{aligned} & \Pr[X_{\hat{t}} = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] \\ & \geq \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ are distinct}] - n\epsilon \\ & \geq 1/e + \delta' \end{aligned}$$

for some $\delta' > 0$, where we have used (1) for the first inequality and $\epsilon = 1/n^2$ 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{t} , 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 [16, Section 2]. \square

3.3 Proof of the Main Structural Lemma

We prove Lemma 3.4 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 on 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 3.5. 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 full version of this paper.

LEMMA 3.6. *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 $\epsilon > 0$, 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 behavior 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. We need the following definition.

Definition 3.7. Consider a stopping rule \mathbf{r} . Let $\epsilon > 0$, $i \in [n]$, and $V \subseteq \mathbb{N}$. Then \mathbf{r} is (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 $r_i(v_1, \dots, v_i) \in [q - \epsilon, q + \epsilon]$.

Note that (ϵ, i) -value-obliviousness for all $i \in [n]$ is equivalent to ϵ -value-obliviousness. In establishing (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 3.8 (RAMSEY [26]). *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 3.4. Consider a stopping rule \mathbf{r} with guarantee α . By Lemma 3.6, it is without loss of generality to assume that \mathbf{r} is order-oblivious. We fix $\epsilon > 0$ 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 $i \in [k-1]$.

Toward the application of Lemma 3.8, we construct a complete k -uniform hypergraph H with vertex set S^{k-1} . Consider any set $\{v_1, \dots, v_k\} \subseteq S^{k-1}$ of cardinality k such that $v_k > \max\{v_1, \dots, v_{k-1}\}$. Note that there exists a unique $u \in \{1, 2, \dots, \lceil 1/(2\varepsilon) \rceil\}$ such that $r_k(v_1, \dots, v_k) \in [(2u-1) \cdot \varepsilon - \varepsilon, (2u-1) \cdot \varepsilon + \varepsilon)$. Color the hyperedge $\{v_1, \dots, v_k\}$ of H with color u .

By Lemma 3.8 with $c = \lceil 1/2\varepsilon \rceil$ and $d = k$, there exists an infinite set of vertices that induces a complete monochromatic sub-hypergraph of H . We define S^k to be such a set inducing a monochromatic sub-hypergraph of H with color u . Now set $q := (2u-1) \cdot \varepsilon$ and consider distinct $v_1, \dots, v_k \in S^k$ with $v_k > \max\{v_1, \dots, v_{k-1}\}$. Since the edge $\{v_1, \dots, v_k\}$ in H has color u , $r_k(v_{\pi(1)}, \dots, v_{\pi(k-1)}, v_k) \in [q - \varepsilon, q + \varepsilon)$ for some permutation $\pi \in \mathcal{S}_{k-1}$. But since \mathbf{r} is order-oblivious, also $r_k(v_1, \dots, v_{k-1}, v_k) \in [q - \varepsilon, q + \varepsilon)$. So \mathbf{r} is (ε, k) -value oblivious on S^k . This completes the induction step and the proof. \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 3.9. *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 4.6 in the full version of this paper.

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_t conditioned on stopping at that random variable is at least $\mathbb{E}[\max\{X_1, \dots, X_n\}]$, since under this condition X_t 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 the full version of this paper, 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*

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

First note that a guarantee of $1 - 1/e - \epsilon$ with $O_\epsilon(n)$ samples follows from a result of Ehsani et al. [13] by observing that $O_\epsilon(n)$ samples provide a sufficiently good approximation to the $1/e$ -quantile of the distribution of $\max\{X_1, \dots, X_n\}$. 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.

Suppose we were given access to $n(n - 1) \in \Theta(n^2)$ samples. Then we could 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 it 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.¹

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

¹An alternative approach for obtaining $n - 1$ fresh-looking samples for each time step i from a single set of $n - 1$ samples $\{S_1, \dots, S_{n-1}\}$ is to select a uniformly random subset of size $n - 1$ from all $n + i - 2$ random draws $\{S_1, \dots, S_{n-1}, X_1, \dots, X_{i-1}\}$ seen so far.

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

recursively as 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 $R_n^i = \max\{R_1^i, \dots, R_{n-1}^i\}$. 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

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

$$= \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_n^t < \max\{R_1^t, \dots, R_n^t\} \right] = \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid \neg \xi_t \right].$$

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. For the third equality, notice that it suffices to show that

$$\Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_n^t < \max\{R_1^t, \dots, R_n^t\} \right] = \Pr \left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg \xi_j \right) \mid R_n^t = \max\{R_1^t, \dots, R_n^t\} \right].$$

So in both cases we evaluate the conditional probability of $\xi_i \cap (\bigcap_{j=t+1}^{i-1} \neg \xi_j)$, i.e., that X_{t+1} to X_{i-1} fail to pass the random thresholds that we set for them and that X_i passes the random threshold that we set for it. First notice that R_1^t, \dots, R_n^t are i.i.d. random variables. The reason is the following: By construction, $\{R_1^t, \dots, R_n^t\}$ is a random subset of $\{S_1, \dots, S_{n-1}, X_1, \dots, X_t\}$, where the choice does not depend on the values of $S_1, \dots, S_{n-1}, X_1, \dots, X_t$. Hence, since $S_1, \dots, S_{n-1}, X_1, \dots, X_t$ are i.i.d., so are R_1^t, \dots, R_n^t . Furthermore note that, for any $t' > t$, all of R_1^t, \dots, R_n^t are among $R_1^{t'}, \dots, R_n^{t'}$ with *equal* probability (as they always remain among the considered variables with equal probability). Hence, for the thresholds we set for X_{t+1}, \dots, X_i , the roles of R_1^t, \dots, R_n^t are indistinguishable, and therefore, for the events at hand, the conditions whether any of R_1^t, \dots, R_n^t is the maximum among them are indistinguishable as well. This establishes the equality and thus 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

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 full version of this paper.

PROPOSITION 4.3. *Let $\epsilon > 0$. 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 ,*

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

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 full version of this paper.

PROPOSITION 4.4. *Let $\epsilon > 0$. 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),*

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

4.4 A Parametric Lower Bound

We generalize our lower bound from Theorem 4.1 to γn samples when $\gamma < 1$. The idea is to reinterpret some amount of values from X_1, \dots, X_n as samples, so that the number of remaining values equals the number of samples and Algorithm 1 can be used.

COROLLARY 4.5. *Let $\gamma \in [0, 1]$ and X_1, X_2, \dots, X_n be i.i.d. random variables from an unknown distribution F . Further assume $\gamma n + n$ to be an even number. Then there exists an $(\gamma n, n)$ -stopping-rule with stopping time τ such that*

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

PROOF. Let $n' := \frac{1 + \gamma}{2} \cdot n \in \mathbb{N}$. Define $S'_i := S_i$ for all $i \in [\gamma n]$, $S'_{\gamma n + i} := X_i$ for all $i \in [n' - \gamma n]$, and $X'_i := X_{n' - \gamma n + i}$ for all $i \in [n']$. Note that $X'_{n'} = X_n$, so this assignment is well-defined. We use Algorithm 1 with associated stopping time τ on $X'_1, \dots, X'_{n'}$, with samples $S'_1, \dots, S'_{n'-1}$. Then by applying Theorem 4.1 we get

$$\mathbb{E}[X'_\tau] \geq \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\max\{X'_1, \dots, X'_{n'}\}] \geq \frac{1 + \gamma}{2} \cdot \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

as claimed. \square

4.5 Close to Tight Upper Bound

While an improvement over the bound of $1 - 1/e \approx 0.632$ remains conceivable 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 full version of this paper.

THEOREM 4.6. *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*




Fig. 2. Visualization of the parametric lower bound (solid) and the parametric upper bound (dashed).

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

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

See Figure 2 for a visualization. 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 [8].

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 [8], if we have access to $O(n^2)$ samples from the distribution. We provide details and the proof in the full version of this paper.

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

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

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

ACKNOWLEDGMENTS

Discussions with Fábio Botler and valuable feedback from the anonymous referees are gratefully acknowledged. This work was funded in part through an Amazon Research Award, the DAAD PRIME program, CONICYT Grant PII 20150140, and the European Research Council under Grant Agreement No. 691672.

REFERENCES

- [1] M. Abolhassani, S. Ehsani, H. Esfandiari, M. Hajiaghayi, R. D. Kleinberg, and B. Lucier. 2017. Beating $1 - 1/e$ for Ordered Prophets. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*. 61–71.
- [2] S. Alaei. 2014. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. *SIAM J. Comput.* 43, 2 (2014), 930–972.
- [3] P. D. Azar, R. Kleinberg, and S. M. Weinberg. 2014. Prophet Inequalities with Limited Information. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1358–1377.
- [4] M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer. 2017. Posting Prices with Unknown Distributions. *ACM Transactions on Economics and Computation* 5, 2 (2017), 13:1–13:20.
- [5] B. A. Berezovskiy and A. V. Gnedin. 1984. *The Best Choice Problem (In Russian)*. Nauka, Moscow.
- [6] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing*. 311–320.
- [7] S. Chawla, J. B. Miller, and Y. Teng. 2019. Posted Pricing for Online Resource Allocation: Beyond Subadditive Values. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1962–1981.
- [8] J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. 2017. Posted Price Mechanisms for a Random Stream of Customers. In *Proceedings of the 18th ACM Conference on Economics and Computation*. 169–186.
- [9] Arnoud V. den Boer. 2015. Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions. *Surveys in Operations Research and Management Science* 20, 1 (2015), 1–18.
- [10] P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. 2017. Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs. In *Proceedings of the 58th Symposium on Foundations of Computer Science*. 540–551.
- [11] P. Dütting and R. D. Kleinberg. 2015. Polymatroid Prophet Inequalities. In *Proceedings of the 23rd European Symposium on Algorithms*. 437–449.
- [12] A. Dvoretzky, J. Kiefer, and J. Wolfowitz. 1956. Asymptotic Minimax Character of the Sample Distribution Function and of the Classical Multinomial Estimator. *Annals of Mathematical Statistics* 27, 3 (1956), 642–669.
- [13] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla. 2018. Prophet Secretary for Combinatorial Auctions and Matroids. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*. 700–714.
- [14] M. Feldman, N. Gravin, and B. Lucier. 2015. Combinatorial Auctions via Posted Prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*. 123–135.
- [15] M. Feldman, O. Svensson, and R. Zenklusen. 2016. Online Contention Resolution Schemes. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1014–1033.
- [16] T. S. Ferguson. 1989. Who Solved the Secretary Problem? *Statist. Sci.* 4, 3 (1989), 282–289.
- [17] J. P. Gilbert and F. Mosteller. 1966. Recognizing the Maximum of a Sequence. *J. Amer. Statist. Assoc.* 61 (1966), 35–73.
- [18] A. Goldenshluger and A. Zeevi. 2017. Optimal Stopping of a Random Sequence with Unknown Distribution. (2017). Working paper.
- [19] M. Hajiaghayi, R. D. Kleinberg, and T. W. Sandholm. 2007. Automated Mechanism Design and Prophet Inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*. 58–65.
- [20] T. Hill and R. P. Kertz. 1992. A Survey of Prophet Inequalities in Optimal Stopping Theory. *Contemp. Math.* 125 (1992), 191–207.
- [21] T. P. Hill and R. P. Kertz. 1982. Comparisons of Stop Rule and Supremum Expectations of I.I.D. Random Variables. *Annals of Probability* 10, 2 (1982), 336–345.
- [22] R. P. Kertz. 1986. Stop Rule and Supremum Expectations of I.I.D. Random Variables: A Complete Comparison by Conjugate Duality. *Journal of Multivariate Analysis* 19 (1986), 88–112.
- [23] R. D. Kleinberg and S. M. Weinberg. 2012. Matroid Prophet Inequalities. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*. 123–136.
- [24] U. Krengel and L. Sucheston. 1977. Semiamarts and Finite Values. *Bull. Amer. Math. Soc.* 83 (1977), 745–747.
- [25] U. Krengel and L. Sucheston. 1978. On Semiamarts, Amarts, and Processes with Finite Value. *Advances in Probability and Related Topics* 4 (1978), 197–266.
- [26] F. P. Ramsey. 1930. On a Problem of Formal Logic. *Proceedings of the London Mathematical Society* 30 (1930), 264–286.
- [27] A. Rubinstein. 2016. Beyond Matroids: Secretary Problem and Prophet Inequality with General Constraints. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*. 324–332.
- [28] A. Rubinstein and S. Singla. 2017. Combinatorial Prophet Inequalities. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1671–1687.
- [29] E. Samuel-Cahn. 1984. Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *Annals of Probability* 12, 4 (1984), 1213–1216.
- [30] J. Wang. 2018. The Prophet Inequality Can Be Solved Optimally with a Single Set of Samples. arXiv:1812.10563.